

Exploring strategies for coupled 4D-Var data assimilation using an idealised atmosphere–ocean model

By POLLY J. SMITH^{1,2*}, ALISON M. FOWLER^{1,2} and AMOS S. LAWLESS^{1,2}, ¹*School of Mathematical and Physical Sciences, University of Reading, Reading, UK*; ²*National Centre for Earth Observation, University of Reading, Reading, UK*

(Manuscript received 17 December 2014; in final form 9 June 2015)

ABSTRACT

Operational forecasting centres are currently developing data assimilation systems for coupled atmosphere–ocean models. Strongly coupled assimilation, in which a single assimilation system is applied to a coupled model, presents significant technical and scientific challenges. Hence weakly coupled assimilation systems are being developed as a first step, in which the coupled model is used to compare the current state estimate with observations, but corrections to the atmosphere and ocean initial conditions are then calculated independently. In this paper, we provide a comprehensive description of the different coupled assimilation methodologies in the context of four-dimensional variational assimilation (4D-Var) and use an idealised framework to assess the expected benefits of moving towards coupled data assimilation. We implement an incremental 4D-Var system within an idealised single-column atmosphere–ocean model. The system has the capability to run both strongly and weakly coupled assimilations as well as uncoupled atmosphere- or ocean-only assimilations, thus allowing a systematic comparison of the different strategies for treating the coupled data assimilation problem. We present results from a series of identical twin experiments devised to investigate the behaviour and sensitivities of the different approaches. Overall, our study demonstrates the potential benefits that may be expected from coupled data assimilation. When compared to uncoupled initialisation, coupled assimilation is able to produce more balanced initial analysis fields, thus reducing initialisation shock and its impact on the subsequent forecast. Single observation experiments demonstrate how coupled assimilation systems are able to pass information between the atmosphere and ocean and therefore use near-surface data to greater effect. We show that much of this benefit may also be gained from a weakly coupled assimilation system, but that this can be sensitive to the parameters used in the assimilation.

Keywords: incremental four-dimensional variational data assimilation, single-column model, KPP mixed layer model, initialisation, strongly coupled, weakly coupled

1. Introduction

The successful application of data assimilation techniques to operational numerical weather prediction and ocean forecasting systems, together with increasing availability of near-surface observations from new satellite missions, has led to an increased interest in their potential for use in the initialisation of coupled atmosphere–ocean models. To produce reliable predictions across seasonal to decadal time scales, we need to simulate the evolution of the atmosphere and ocean coupled together. Coupled models have been used operationally for seasonal and longer-range forecasting for

a number of years, and are now being considered for shorter term prediction. Typically, the initial conditions for these forecasts are provided by combining analyses from independent (uncoupled) ocean and atmosphere assimilation systems (Balmaseda and Anderson, 2009). This approach ignores interactions between the systems and this inconsistency can cause imbalance such that the initial conditions are far from the natural state of the coupled system. When the coupled forecast is initialised, the model adjusts itself towards its preferred climatology; this adjustment can produce rapid shocks at the air–sea interface during the early stages of the forecast, a process referred to as initialisation shock (Balmaseda, 2012). It also means that near-surface data are not fully utilised.

The development of coupled atmosphere–ocean data assimilation systems presents a number of scientific and

*Corresponding author.
email: p.j.smith@reading.ac.uk

technical challenges (Murphy et al., 2010; Lawless, 2012) and requires a significant amount of resources to be made possible operationally. Yet such systems offer a long list of potential benefits including improved use of near-surface observations, reduction of initialisation shocks in coupled forecasts, and generation of a consistent system state for the initialisation of coupled forecasts across all timescales. In addition, coupled reanalyses offer the potential for greater understanding and representation of air–sea exchange processes, in turn facilitating more accurate prediction of phenomena such as El Niño and the Madden–Julian Oscillation (MJO) in which air–sea interaction plays an important role.

Studies have shown that, in certain regions, the initialisation of coupled models can enhance the skill of decadal predictions for the first five or so years of the forecast [Meehl et al. (2014) and references therein]. Although it is widely accepted that coupled data assimilation has a central role in improving our ability to generate consistent and accurate initial conditions for coupled atmosphere–ocean forecasting, it is still a relatively young area of research. Hence there has so far only been a limited amount of work in this field. An assortment of strategies for using observed data to improve coupled model initialisation have been explored with varying degrees of success; these include sea surface temperature (SST) nudging or relaxation (e.g. Keenlyside et al., 2008), anomaly initialisation/bias-blind assimilation (e.g. Pierce et al., 2004), anomaly coupling (Pohlmann et al., 2009), and variants of the full uncoupled initialisation approach (Balmaseda and Anderson, 2009). Work has mainly been focussed on improving ocean initial conditions with a lack of fully consistent treatment of air–sea feedback mechanisms.

There are groups exploring more comprehensive approaches that aim to produce more dynamically balanced initial ocean–atmosphere states. The Japan Agency for Marine–Earth Science and Technology (JAMSTEC) are working towards a coupled 4D-Var data assimilation system for their Coupled model for the Earth Simulator (CFES), a fully coupled global climate model. Sugiura et al. (2008) describes the development of a first step 4D-Var system for estimating ocean initial conditions together with adjustment parameters of the bulk flux formulae. Their approach is focussed on representing slow time scales only, filtering out fast atmospheric modes by using 10-d mean states. Whilst this enables them to better represent several key seasonal to interannual climate events in the tropical Pacific and Indian Ocean region, including the El Niño, it would not be suitable for atmospheric reanalyses or for initialising medium-range forecasts.

So far, most coupled data assimilation work in the published literature has employed ensemble-based assimilation methods rather than variational-based assimilation methods.

Tardif et al. (2014) use an idealised low-dimension atmosphere–ocean climate model with an Ensemble Kalman Filter (EnKF) to explore strategies for ensemble coupled data assimilation. Their model represents an idealisation of the midlatitude North Atlantic climate system and is designed to allow experiments on very long timescales in order to assess the EnKF approach in terms of effectiveness for initialisation of the meridional overturning circulation (MOC). In twin experiments using 50-yr windows from a 5000-yr reference simulation, they found that forcing the idealised ocean model with atmospheric analyses was inefficient at recovering the MOC due to slow convergence of the solutions. In contrast, coupled assimilation produced accurate MOC analyses, even when only atmospheric observations were assimilated.

In a larger scale study, Zhang et al. (2007) describe a coupled assimilation system consisting of an EnKF applied to the National Oceanic and Atmospheric Administration (NOAA) Geophysical Fluid Dynamics Laboratory (GFDL) global fully coupled climate model for the initialisation of seasonal and decadal forecasts. The system is evaluated in a series of twin experiments assimilating atmosphere-only or ocean-only observations, but not both together. Although the system shows good skill in reconstructing seasonal and decadal ocean variability and trends, it fails to fully realise the potential benefit of surface and near-surface observational data.

In this paper, we explore some of the fundamental questions in the design of coupled variational data assimilation systems within the context of an idealised one-dimensional (1D) column coupled atmosphere–ocean model. The system is designed to enable the effective exploration of various approaches to performing coupled model data assimilation whilst avoiding many of the issues associated with more complex models and allows us to perform experiments that would not be feasible in operational scale systems. We employ an incremental four-dimensional variational data assimilation (4D-Var) scheme (Courtier et al., 1994; Lawless et al., 2005; Lawless, 2013) to reflect the coupled assimilation systems currently being developed at the European Centre for Medium Range Weather Forecasts (ECMWF) and UK Met Office. The problem of variational data assimilation is to find the initial state such that the model forecast best fits the available observations over a given time window, subject to the state remaining close to a given *a priori*, or background, estimate and allowing for the errors in each. This best estimate is known as the *analysis* and should be consistent with both the observations and the system dynamics. The standard 4D-Var problem is formulated as the minimisation of a non-linear weighted least squares cost function; in the incremental approach, the non-linear problem is instead approximated by a sequence of linear least squares problems.

Rather than searching for the initial state directly, we solve in terms of increments with respect to an initial background state; this is done iteratively in a series of linearised inner-loop cost function minimisations and non-linear outer-loop update steps.

Strongly or fully coupled variational assimilation treats the atmosphere and ocean as a single coherent system, using the coupled model in both the inner- and outer-loops. This approach is able to pass information between the atmosphere and ocean, and therefore enables observations of atmospheric variables to influence the ocean increments and vice versa. This is expected to lead to better use of near-surface observations, such as scatterometer winds and SST, that depend on both the atmosphere and ocean state, and to produce a more physically balanced analysis. Although there are currently no plans to move towards strongly coupled systems at operational centres, this approach represents the quintessential coupled assimilation system and implementing it in our idealised system allows us to better assess the potential of intermediate, or *weakly coupled*, approaches.

As a first step towards the implementation of operational coupled data assimilation, centres such as the ECMWF and UK Met Office are developing prototype weakly coupled assimilation systems (Laloyaux et al., 2014, 2015; Brassington et al., 2015; Lea et al., 2015). Weakly coupled incremental 4D-Var makes use of the incremental inner- and outer-loop structure; the coupled model is used to provide the initial atmosphere and ocean background states and non-linear trajectory for separate (uncoupled) inner-loop atmosphere and ocean minimisations; the uncoupled analysis increments are then fed back into the coupled model for the next outer-loop forecast. Unlike strongly coupled assimilation, the weakly coupled approach does not allow for cross-covariance information between the atmosphere and ocean. This means that the atmosphere (ocean) observations cannot affect the ocean (atmosphere) analyses unless multiple outer-loops are performed.

The purpose of this paper is to (1) provide a comprehensive description of the different coupled 4D-Var data assimilation methodologies, and (2) use our idealised framework to assess the benefits expected in moving towards coupled data assimilation systems. Although the greatest benefits are anticipated to be attained with strongly coupled assimilation, we investigate whether the weakly coupled approaches being pursued by operational centres are likely to provide a determinable improvement on the current uncoupled systems. We consider if the potential added benefits of strongly coupled assimilation ultimately outweigh the challenges their development presents.

We begin, in Section 2, with the formulation of the general incremental 4D-Var algorithm and a description of the different approaches to coupled atmosphere–ocean

4D-Var data assimilation. We introduce our coupled 1D model system in Section 3. In Section 4, we give details of a set of identical twin experiments designed to investigate and compare the behaviour and sensitivities of the different approaches. Results are presented in Section 5. A summary and conclusions are given in Section 6.

2. Incremental 4D-Var data assimilation

Variational methods form the basis of most operational numerical weather prediction (NWP) data assimilation systems (Gauthier et al., 1999, 2007; Rabier et al., 2000; Rawlins et al., 2007; Huang et al., 2009). Our system has therefore been designed using the incremental 4D-Var approach. In this formulation, the solution to the full non-linear 4D-Var minimisation problem is replaced by a sequence of minimisations of linear quadratic cost functions such that the control variable in the minimisation problem is the increment to the current estimate rather than the model state itself. The method was originally developed to overcome the cost and practical difficulties involved in solving the complete non-linear problem directly in large scale systems (Courtier et al., 1994). We choose to employ the incremental 4D-Var formulation for this study as it allows us not only to emulate the methodologies being developed for operational systems, but also to explore the type of benefits that could be gained by moving towards strongly coupled assimilation systems, thereby providing a benchmark for the assessment of weakly coupled assimilation systems. We describe each of the different coupled 4D-Var assimilation strategies in detail in Section 2.1. To aid these descriptions, we begin with an outline of the steps of the general incremental 4D-Var algorithm.

Let $\mathbf{x}_i \in \mathbb{R}^m$ denote the model state vector, representing the system state at a given time t_i . Then given the discrete non-linear dynamical system model

$$\mathbf{x}_i = \mathcal{M}(t_i, t_0, \mathbf{x}_0), \quad i = 0, \dots, n, \quad (1)$$

a background, or first guess, $\mathbf{x}_0^b \in \mathbb{R}^m$, at t_0 , and imperfect observations $\mathbf{y}_i \in \mathbb{R}^{r_i}$ at times t_i , $i = 0, \dots, n$, the full non-linear 4D-Var problem is to find the initial model state, \mathbf{x}_0 , that minimises the cost function

$$J(\mathbf{x}_0) = \frac{1}{2} (\mathbf{x}_0^b - \mathbf{x}_0)^T \mathbf{B}_0^{-1} (\mathbf{x}_0^b - \mathbf{x}_0) + \frac{1}{2} \sum_{i=0}^n (\mathbf{y}_i - h_i(\mathbf{x}_i))^T \mathbf{R}_i^{-1} (\mathbf{y}_i - h_i(\mathbf{x}_i)). \quad (2)$$

Here $h_i : \mathbb{R}^m \rightarrow \mathbb{R}^{r_i}$ is a (generally) non-linear observation operator and $\mathbf{B}_0 \in \mathbb{R}^{m \times m}$ and $\mathbf{R}_i \in \mathbb{R}^{r_i \times r_i}$ are the background and observation error covariance matrices respectively.

For the incremental 4D-Var approach, rather than minimising [eq. (2)] directly we define the increment, at a given time t_i and outer-loop iteration k , as

$$\delta \mathbf{x}_i^{(k)} = \mathbf{x}_i^{(k+1)} - \mathbf{x}_i^{(k)}, \quad (3)$$

and solve iteratively as described below (Lawless et al., 2005).

For $k=0, 1, \dots, K$ outer-loops, or until desired convergence is reached:

- (i) For the first iteration set $\mathbf{x}_0^{(0)} = \mathbf{x}_0^b$.
- (ii) Run the non-linear model [eq. (1)] to obtain $\mathbf{x}_i^{(k)}$ at each time t_i .
- (iii) Compute the innovations

$$\mathbf{d}_i^{(k)} = \mathbf{y}_i - h_i(\mathbf{x}_i^{(k)}). \quad (4)$$

- (iv) Minimise the least squares cost function

$$\begin{aligned} J^{(k)}(\delta \mathbf{x}_0^{(k)}) &= \frac{1}{2} \left((\mathbf{x}_0^b - \mathbf{x}_0^{(k)}) - \delta \mathbf{x}_0^{(k)} \right)^T \mathbf{B}_0^{-1} \left((\mathbf{x}_0^b - \mathbf{x}_0^{(k)}) - \delta \mathbf{x}_0^{(k)} \right) \\ &+ \frac{1}{2} \sum_{i=0}^n \left(\mathbf{d}_i^{(k)} - \mathbf{H}_i \delta \mathbf{x}_i^{(k)} \right)^T \mathbf{R}_i^{-1} \left(\mathbf{d}_i^{(k)} - \mathbf{H}_i \delta \mathbf{x}_i^{(k)} \right), \\ &:= J_b^{(k)} + J_o^{(k)}, \end{aligned} \quad (5)$$

subject to

$$\delta \mathbf{x}_i^{(k)} = \mathbf{M}(t_i, t_0, \mathbf{x}^{(k)}) \delta \mathbf{x}_0^{(k)}. \quad (6)$$

- (v) Update $\mathbf{x}_0^{(k+1)} = \mathbf{x}_0^{(k)} + \delta \mathbf{x}_0^{(k)}$, and return to step (ii).

In eq. (5), the operator $\mathbf{H}_i \in \mathbb{R}^{r_i \times m}$ is the tangent linear (TL) of the non-linear observation operator, h_i , and \mathbf{M} is the TL of the non-linear model operator \mathcal{M} . For each outer-loop, k , the linearised operators \mathbf{H} and \mathbf{M} are evaluated at the current estimate of the non-linear trajectory, $\mathbf{x}^{(k)}$, referred to as the linearisation state.

Step (iv) is referred to as the ‘inner-loop’. The minimisation of the cost function [eq. (5)] is performed iteratively using a gradient descent algorithm. For each iteration of the inner-loop minimisation, the TL model [eq. (6)] is integrated to give the evolution of the increment for the cost function computation [eq. (5)] and the adjoint of the TL model, \mathbf{M}^T , is integrated to obtain the cost function gradient. In this study, we employ an off the shelf optimisation algorithm based on the conjugate gradient method (Shanno, 1978; Shanno and Phua, 1980).

2.1. Coupled data assimilation

Our system has been designed to enable several different 4D-Var configurations: an uncoupled atmosphere-only or ocean-only assimilation, a weakly coupled assimilation and a strongly coupled assimilation, thus allowing a systematic comparison of the different strategies for treating the

coupled 4D-Var data assimilation problem. In this section, we give details of each algorithm and highlight the main differences between them.

2.1.1. Strongly coupled incremental 4D-Var. For the strongly (or fully) coupled assimilation system, the state vector, \mathbf{x} , and the incremental 4D-Var control vector, $\delta \mathbf{x}$, consist of both the atmosphere and ocean prognostic variables. The coupled model is used in both the outer and inner-loops; the non-linear coupled model is used in the outer-loops to generate the linearisation trajectory [eq. (1)] and compute the innovation vectors [eq. (4)], and the inner-loop cost function minimisation is performed using the TL and adjoint of the coupled non-linear model in step (iv). Information is exchanged between the atmosphere and ocean components at regular, specified time intervals; the SST from the ocean model is used in the computation of the atmospheric lower boundary conditions, and the surface heat, moisture and momentum fluxes from the atmosphere model provide the ocean surface boundary conditions.

The incremental 4D-Var algorithm implicitly evolves the background error covariances across the assimilation window according to the TL model dynamics (e.g. Thépaut et al., 1993, 1996). This acts to modify the prior background error variance estimates and induce non-zero correlations between model variables. The use of the fully coupled TL and adjoint models in the inner-loops of the strongly coupled assimilation system means that we expect cross-covariance information to be generated between the atmosphere and ocean fields. This allows observations of one fluid to produce analysis increments in the other and is therefore expected to generate more consistent analyses. The design of this system also has the advantage of allowing for cross-covariances between the atmosphere and ocean errors to be explicitly prescribed *a priori*. We note, however, that this is a non-trivial matter; research on how to characterise and represent atmosphere–ocean cross-covariances within coupled data assimilation systems is currently underway and will be addressed in a future paper.

2.1.2. Uncoupled incremental 4D-Var. The uncoupled atmosphere and ocean assimilation systems are completely independent. Here, the state and incremental 4D-Var control vectors are comprised of the atmosphere or ocean prognostic variables only and a separate inner-loop cost function is used for each model. For the atmosphere (ocean), the outer-loop linearisation trajectory [eq. (1)] is taken from a run of the atmosphere-only (ocean-only) non-linear model, the innovation vectors [step (iii)] are computed using the available atmosphere (ocean) observations and the inner-loop minimisation [step (iv)] uses the corresponding

uncoupled atmosphere (ocean) TL and adjoint model. There is no exchange of information between the two systems at any stage; the SST used at the atmosphere bottom boundary and the momentum, heat and freshwater fluxes at the ocean surface boundary are prescribed. Although this approach has its advantages, such as ease of implementation and modularity, it does not allow for cross-covariances between the atmosphere and ocean fields and atmospheric (ocean) observations cannot influence the ocean (atmosphere) analysis. The lack of feedback means that the atmosphere and ocean analysis states are unlikely to be in balance and this can have a negative impact if they are used to initialise a coupled model forecast (Balmaseda and Anderson, 2009).

2.1.3. Weakly coupled incremental 4D-Var. The weakly coupled 4D-Var algorithm is a combination of the strongly and uncoupled algorithms; the system uses a coupled full state vector but uncoupled atmosphere and ocean incremental 4D-Var control vectors. This approach has the advantage that it limits the amount of new technical development required when independent atmosphere and ocean assimilation systems are already in place. The outer-loop linearisation trajectory [eq. (1)] is generated using the coupled non-linear model but separate inner-loop cost functions [step (iv)] are defined for the atmosphere and ocean, using the respective uncoupled atmosphere- or ocean-only TL and adjoint models and assimilating the atmosphere or ocean observations only. Although the computation of the innovations [step (iii)] uses only the atmosphere or ocean observations, the observation-model fit is measured against the coupled model state. The ocean SST from the coupled outer-loop linearisation trajectory is used in the computation of the bottom boundary conditions for the uncoupled atmospheric TL model and the surface heat, moisture and momentum fluxes from the coupled outer-loop linearisation trajectory are used in the computation of the surface boundary conditions for the uncoupled ocean TL model. Once the uncoupled atmosphere and ocean inner-loop minimisations have been performed, the uncoupled atmosphere and ocean analysis increments are combined and added to the current guess to provide the initial coupled state for the next outer-loop iteration.

Analogous to the uncoupled case, the separation of the atmosphere and ocean TL model components in the inner-loops of the weakly coupled system means that cross-covariances between the atmosphere and ocean are ignored; they can only be generated between atmosphere fields or between ocean fields. However, as we demonstrate in Sections 5.4 and 5.5, observations of one fluid are able to influence the analysis of the other if multiple outer-loops are performed due to the linearisation state being updated.

Here, we have presented the coupled 4D-Var assimilation strategies in their cleanest forms. It should be noted that in practice there will be variations in their application. For example, the uncoupled analysis systems at the ECMWF and UK Met Office run the atmosphere component with a prescribed SST but then use the updated fluxes from this analysis to constrain the ocean assimilation system. It is also not necessary for the uncoupled and weakly coupled systems to use the same assimilation window length for the atmosphere and ocean, or even the same assimilation scheme. Both the existing uncoupled analysis systems and the weakly coupled systems currently under development at the ECMWF and UK Met Office use 4D-Var for the atmosphere and 3D-Var FGAT (first guess at appropriate time) for the ocean (Laloyaux et al., 2014, 2015; Lea et al., 2015).

3. The coupled model

The objective of this study is to gain a greater theoretical understanding of the coupled atmosphere–ocean data assimilation problem by exploring and comparing the behaviours of the coupling strategies presented in Section 2.1. Idealised models offer an effective framework for investigating and advancing new methods, avoiding unnecessary complexities that can obscure results. Using a simplified system allows us to perform a range and quantity of experiments that would require a significant amount of technical development and resources to execute in a full scale system.

In many cases, finding balanced solutions to the coupled atmosphere–ocean assimilation problem is primarily a vertical problem of the two boundary layers. A 1D column atmosphere–ocean model framework therefore offers a tractable and relevant approach. Whilst it is preferable to keep the model as simple as possible from a developmental point of view, it is important to ensure that processes crucial to realistic air–sea interaction, such as the diurnal SST cycle and evolution of surface forcing, are adequately represented. Our new system has been built by coupling the ECMWF single-column atmospheric model (SCM) to a single-column K-Profile Parameterisation (KPP) ocean mixed layer model. The use of these models ensures that the simplified system retains the key elements of coupling processes in a fully coupled ocean–atmosphere model without being overly complex.

3.1. Non-linear models

The atmosphere model solves the primitive equations for temperature, T , specific humidity, q , and zonal, u , and meridional, v , wind components, formulated in non-spherical co-ordinates (Simmons and Burridge, 1981; Ritchie et al., 1995). Compared to the original SCM, our model does not

include the parameterisation of physical processes such as radiation, subgrid-scale orographic drag, convection, clouds and surface/soil processes; we also use a simplified vertical diffusion scheme. For the ocean, the evolution of the temperature, θ , salinity, s , and zonal and meridional currents u_o , v_o are described following the formulation of (Large et al., 1994). Further details, including the model equations, are given in the Appendix.

The coupling of the two models takes place at the atmosphere–ocean boundary. In the atmosphere, the lower boundary conditions for temperature and specific humidity depend on the SST (temperature at top level of ocean model which corresponds to a depth of 1 m) and saturation specific humidity, q_{sat} (SST), and a no-slip condition is used for the u and v wind components. The ocean surface boundary conditions for temperature and salinity are given by the surface kinematic fluxes of heat and salt which in turn depend on the net long wave radiation and latent and sensible heat fluxes. The surface boundary conditions for the u and v components of the current are given by the zonal and meridional components of the kinematic momentum flux. At each model time step, the atmosphere model computes and passes the latent and sensible heat fluxes and the horizontal components of the surface momentum flux to the ocean model. The updated ocean model SST is passed back to the atmosphere where it is used in the computation of the atmosphere lower boundary conditions for the next step.

A fuller description of the individual atmosphere and ocean non-linear model components and their coupling is given in the Appendix. This system combined with our 4D-Var schemes provides a unique and tractable framework for addressing the coupled atmosphere–ocean assimilation problem.

As part of the assimilation system development, the simplified non-linear model was validated against the original (full physics) version of the ECMWF SCM code. As expected, we see small differences in the evolution of both the prognostic variables and surface fluxes but in general we find that using the simplified physics provides a good approximation to the full physics in the coupled model. Where there are differences the simplified model still produces an evolution that is physically reasonable, with a diurnal cycle in the ocean SST and mixed layer depth and appropriate atmosphere–ocean fluxes; see, for example, the truth trajectory (black line) in Fig. 4. We are therefore confident that the model is sufficient for assessing the different assimilation strategies.

3.2. TL and adjoint models

In order to be able to compute the cost function and its gradient for each inner-loop we need to develop the tangent

linear (TL) and adjoint models. A particular issue worth noting is the linearisation of the atmosphere and ocean vertical turbulent flux parameterisations. The formulation of the diffusion coefficients K_ϕ in both the atmosphere and ocean vertical diffusion schemes [eqs. (A6) and (A13)] is strongly non-linear and its linearisation has been shown to be unstable (Laroche et al., 2002). The simplest way to avoid difficulties associated with this linearisation is to neglect the perturbation of the K_ϕ coefficients. Studies such as Janisková et al. (1999), Mahfouf (1999) and Laroche et al. (2002) have shown that a TL diffusion scheme can still produce reasonable and useful behaviour under this assumption and this approach has been widely adopted in both atmosphere and ocean assimilation systems (e.g. Mahfouf, 1999; Weaver et al., 2003). We are therefore satisfied that this simplification is appropriate for our system. During the assimilation, the K_ϕ are computed for each non-linear outer-loop and then held constant for the inner-loop minimisation.

Although this means we are using an approximate TL model rather than the exact TL, since the adjoint model is derived from the approximate TL model, the inner-loop cost function gradient calculation contains the correct information for convergence of the minimisation problem. The correctness of the tangent linear and adjoint model codes, and the gradient calculation were all verified using standard tests (e.g. Navon et al., 1992; Lawless, 2013).

4. Experimental design

We compare the performance of the strongly coupled, weakly coupled and uncoupled 4D-Var systems via a series of identical twin experiments in which the coupled non-linear model is used to forecast a reference or ‘truth’ trajectory from which synthetic observations are generated. The assimilation systems are then assessed on how well they approximate the initial reference state and subsequent forecast.

4.1. Initial conditions and forcing

The atmospheric initial conditions, surface pressures, and SST data used to force the uncoupled atmosphere system are taken from the ERA Interim Re-analysis¹ (Dee et al., 2011). Fields are available at 6-hourly intervals and can be extracted on model levels so that they do not need pre-processing. These data are also used to estimate the geostrophic wind components, u_g , v_g in eqs. (A1) and (A2), and large scale horizontal forcing terms (see Appendix)

¹ERA Interim Re-analysis data can be downloaded via the ECMWF data server at www.ecmwf.int.

using simple centred finite difference approximations across adjacent latitude and longitude points.

Initial ocean fields are produced by interpolating Mercator Ocean reanalysis data² (Lellouche et al., 2013), onto the KPP model grid. The surface short and long wave radiation forcing fields are computed by running the full physics version of the coupled single-column model and taking 6-hourly snapshots of the diagnostic clear-sky radiation flux fields that are computed as part of the radiation scheme (ECMWF IFS documentation, 2001–2013b). The geostrophic components of the ocean currents (Section A.2.1) are estimated by computing a 10-d rolling average of the Mercator Ocean currents. The surface heat, moisture and momentum fluxes required for uncoupled ocean model integrations are taken from the ERA Interim Re-analysis.

For the experiments presented here, the true initial state, \mathbf{x}_0 , is a 24-hour coupled model forecast from 00:00 UTC on 2nd June 2013 to 00:00 UTC on 3rd June 2013 for the point (188.75°E, 25°N) which is located in the north west Pacific ocean. This forecast was initialised using ERA interim and Mercator Ocean reanalysis data; we denote this initial forecast state as \mathbf{x}_{-24} . We run a forecast rather than initialising from these data directly in order to generate an initial state that is consistent with the coupled model dynamics.

4.2. Background

The initial background state, \mathbf{x}_0^b , is generated by running a second 24-hour coupled model forecast from 00:00 UTC on 2nd June 2013 with perturbed initial data; this data, denoted $\hat{\mathbf{x}}_{-24}$, is generated by adding random Gaussian noise to \mathbf{x}_{-24} ,

$$\hat{\mathbf{x}}_{-24} = \mathbf{x}_{-24} + \boldsymbol{\sigma} \circ \delta\mathbf{x}. \quad (7)$$

Here, the $\delta\mathbf{x}$ are normally distributed, random perturbations and $\boldsymbol{\sigma} \in \mathbb{R}^m$ is a vector of standard deviations; these are computed as the sample standard deviation of the unperturbed coupled model forecast states at each model time step from \mathbf{x}_{-24} to \mathbf{x}_0 . The initial background guess for the assimilation is then given by the perturbed coupled model forecast state after 24 hours, i.e. $\mathbf{x}_0^b = \hat{\mathbf{x}}_0$. The true and initial background states are shown in Fig. 1.

For the purposes of this study, we assume that the background error covariance matrix \mathbf{B}_0 is diagonal, that is, the initial background errors are univariate and spatially uncorrelated; the diagonal elements, σ_b^2 , representing the error variances, are assumed to vary for each model field

and vertical level and are taken to be the squared values of the standard deviations $\boldsymbol{\sigma}$ from eq. (7). The initial background error standard deviation profiles are shown in Fig. 2. The relative magnitude of the standard deviations of the atmosphere fields is greater than the ocean due to the faster timescales. For the ocean temperature and salinity, the variability, and thus the prescribed background error standard deviation, is largest in the turbulent mixed layer region (\sim top 50 m) where timescales are shortest. Moving deeper into the ocean the timescales become longer and the standard deviations become very small.

The assumption of a diagonal matrix \mathbf{B} is a great simplification but is used here as an aid to understanding the implicit evolution of the error covariances by the 4D-Var algorithm. Although we assume that the prior atmosphere and ocean fields are uncorrelated, the incremental 4D-Var algorithm implicitly propagates the background error covariances across the assimilation window according to the TL model dynamics [see Bannister (2008a) and references therein]. This acts to modify the prior background error variance estimates and induce non-zero correlations between model variables.

A simple preconditioning of the inner-loop cost function using the square root of the background error covariance matrix was found to be beneficial in terms of improving the conditioning of the system and allowing convergence of the inner-loop minimisation within a reasonable number of iterations (Courtier, 1997). Preconditioning is common in most operational variational assimilation systems and is often implemented using a control variable transform (Bannister, 2008b).

4.3. Observations

We assume that the model state variables are observed directly to avoid the additional complexity of a non-linear observation operator. Observations are generated by adding uncorrelated random Gaussian errors, with given standard deviations (see Table 1) to the reference trajectory at constant time and space intervals. Observations of atmospheric temperature, and u and v wind components are assimilated at 17 of the 60 atmosphere model levels; these are chosen to approximately correspond to the standard pressure levels used, for example, to report radiosonde data (see Table 2). Observations of ocean temperature, salinity, and zonal and meridional currents are assimilated at 23 of the 35 ocean model levels giving vertical frequency comparable to a XBT profile (see Table 3). Note that since the atmospheric model does not include the parameterisation of processes such as moist convection, clouds and precipitation we do not assimilate observations of specific humidity, q . Ocean observational data are typically available less frequently than atmospheric observational

²Mercator Ocean re-analysis data are available via the MyOcean project Web Portal at www.myocean.eu.org.

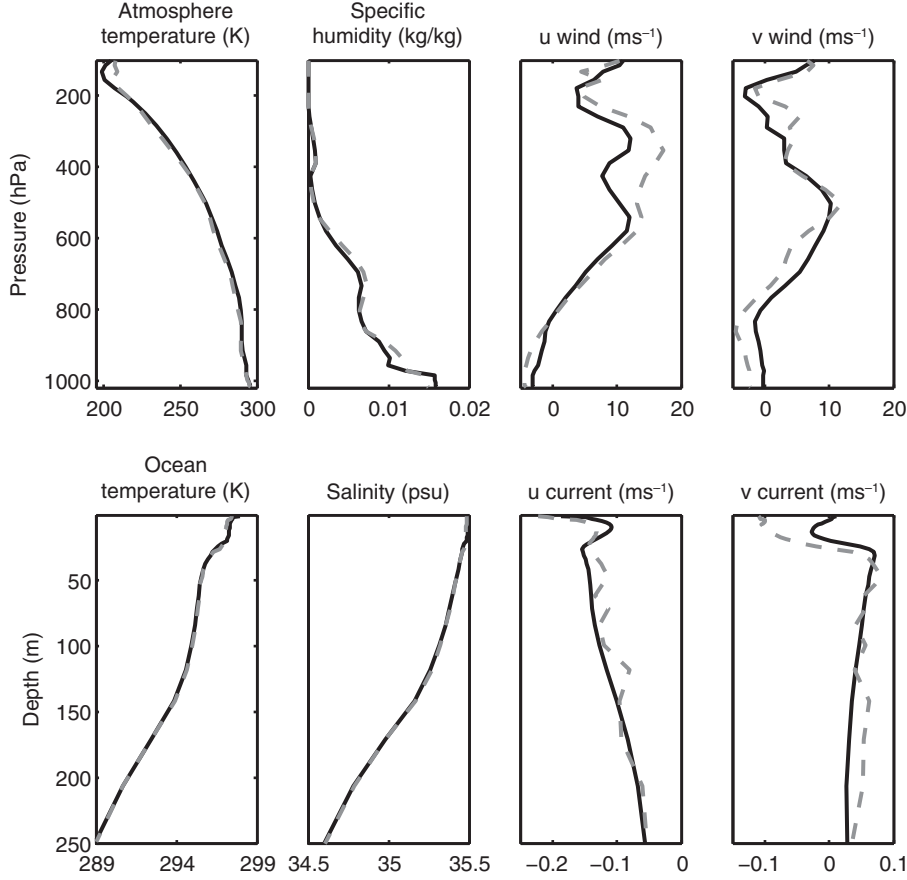


Fig. 1. True (solid black line) and background (dashed grey line) initial states.

data, particularly for certain operational observing systems. The atmosphere and ocean observation frequencies used in our assimilation experiments were chosen to reflect this disparity. Unless otherwise stated, results refer to experiments run with atmosphere observations at 3, 6, 9 and 12 hours, and ocean observations at 6 and 12 hours.

Although it is generally accepted that observation error covariances exist, it is typical to ignore them and in practice it is assumed that the errors in the observational data are spatially and temporally uncorrelated so that the observation error covariance matrices \mathbf{R}_i are diagonal (Daley, 1991). We follow the same approach here but also keep the observation network fixed for the duration of each experiment so that the number of observations $r_i = r$ and $\mathbf{R}_i = \mathbf{R}$ for all i . The observation error variances, σ_o^2 , are assumed to be constant across all vertical levels for each observation type.

5. Assimilation results

Since our aim is to examine the impact of coupled assimilation on the atmosphere and ocean boundary layers we limit

our discussion to this region and focus on the results in the bottom ~ 200 hPa of the atmosphere model (~ 15 levels), and top 50 m (26 levels) of the ocean model. We use a 12-hour assimilation window with three outer-loops and a model time step of 15 minutes. A 12-hour window length is common for atmospheric data assimilation systems, such as the ECMWF IFS. Initial experiments using a greater number of outer-loops showed that the cost function usually converged after around three loops. The inner-loop minimisation is terminated when the relative change in gradient is less than 0.001 (Lawless and Nichols, 2006); for our system, this is typically after around 10–20 iterations.

Figure 3 shows the absolute (truth-analysis) error profiles at initial time t_0 for each of the prognostic model variables for the three assimilation systems. The differences between the analyses are most pronounced in the upper ocean temperature and u_o , v_o current fields. The atmospheric temperature and specific humidity analysis errors are very similar to the initial background errors for all three systems. For specific humidity, this is expected since we do not observe this field. For atmospheric temperature, this may be in part due to the fact that the initial background

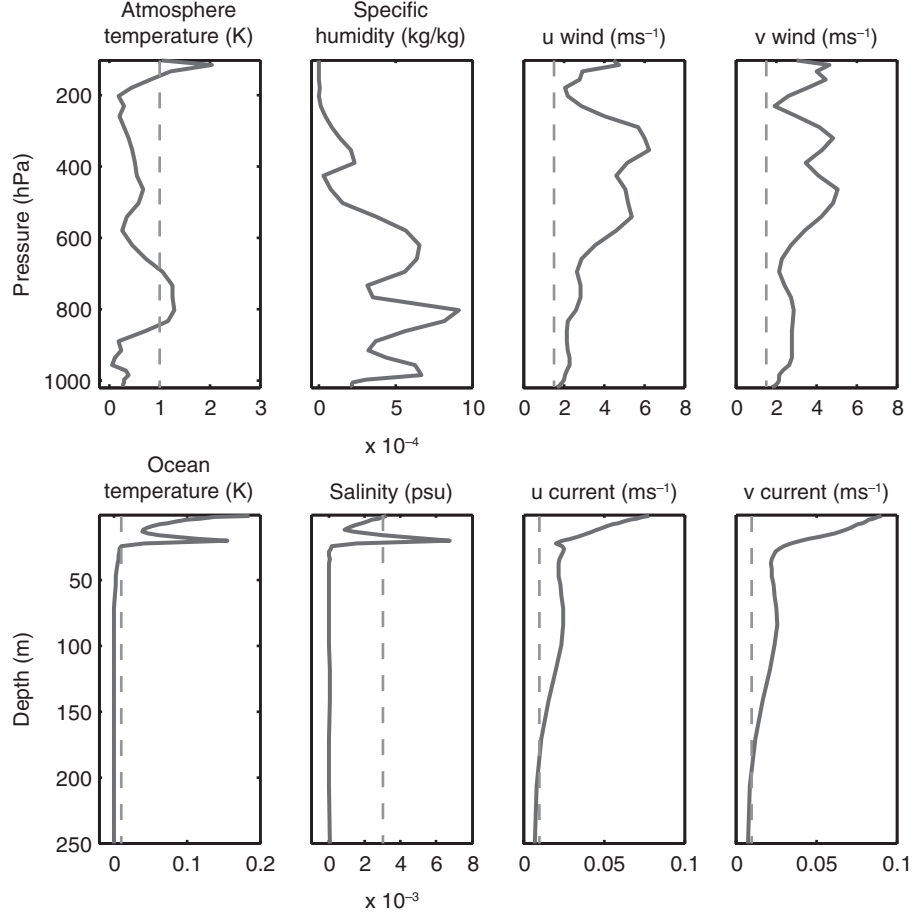


Fig. 2. Initial background (solid grey line) and observation (vertical dashed grey line) error standard deviations.

errors are small in this region. There are also relatively fewer observations in the lower atmosphere compared to the upper ocean. There are clearer improvements in the near-surface u and v wind fields. Here, the analysis errors for the uncoupled and weakly coupled systems are very alike whereas the strongly coupled system appears to use the near-surface ocean current observations to further correct the u and v wind analyses. This is, in part, due to the way the coupled and uncoupled non-linear models and assimilation systems are formulated.

A notable aspect in the ocean analysis errors is at approximately 20 m, which coincides with the mixed layer depth. The mixed layer depth is characterised by a sharp gradient in the temperature and salinity profiles. In the

background estimate, the position of this feature is incorrect. When the assimilation of observations attempts to correct this positional error, instead of shifting the profiles, it erroneously changes the structure of the temperature and salinity profiles so that the error in the analysis is actually increased compared to the background. This is an issue for all three coupling strategies and is a well-documented problem in the atmosphere when assimilating observations of the analogous boundary layer capping inversion (Fowler et al., 2012).

It is not possible to draw conclusions on the performance of each approach from the analysis errors alone. In particular, these results do not give any indication of whether the initial atmosphere and ocean analysis states are in balance.

Table 1. Observation error standard deviations by field

Atmosphere temperature (K)	u wind (ms^{-1})	v wind (ms^{-1})	Ocean temperature (K)	Salinity (psu)	u_o current (ms^{-1})	v_o current (ms^{-1})
1.0	1.5	1.5	0.01	0.003	0.01	0.01

Table 2. Atmosphere observation locations

Model level	Standard pressure level (hPa)	Model full-level pressure value (hPa) ^a
14	10	9.893
17	20	18.815
19	30	28.882
22	50	54.624
23	70	66.623
25	100	95.980
28	150	154.038
30	200	202.230
32	250	257.685
33	300	288.093
36	400	389.233
39	500	501.637
44	700	694.696
49	850	861.497
52	925	935.065
56	1000	995.055
60	n/a	1017.293

^aValues based on a surface pressure value of 1018.5 hPa; model full-level pressure values vary with surface pressure. These levels have been chosen to approximately correspond to the standard pressure levels (hPa).

Since one of the key drivers behind the development of coupled data assimilation systems is generation of a consistent system state for the initialisation of coupled model forecasts, we use the analysis fields at the beginning of the assimilation window to initialise a series of coupled model forecasts; the results are discussed in Sections 5.1–5.3.

5.1. Initialisation shock

A major problem with using analysis states from uncoupled assimilation systems to initialise a coupled model forecast is that the atmosphere and ocean fields may not be balanced and this can lead to initialisation shock. If the initial conditions are not on the coupled model attractor (in these twin experiments also the true attractor), the forecast will experience an adjustment process. In some cases, the adjustment towards the model attractor solution occurs asymptotically but in others it manifests itself as a rapid change in the model fields in the early stages of the forecast (Balmaseda, 2012). The skill of a coupled model forecast depends strongly on the way it is initialised, thus the reduction or elimination of initialisation shock is particularly important in seasonal forecasting (Balmaseda and Anderson, 2009).

Figure 4 compares the SST and surface fluxes for the first 48 hours of each coupled model forecast against the truth trajectory and also a forecast initialised from the initial background state (i.e. no assimilation). In all cases, the forecast eventually tracks the true trajectory fairly well

Table 3. Ocean observation locations

Model level	depth (m)
1	1.000
3	3.069
5	5.277
8	8.848
10	11.406
13	15.538
16	20.173
18	23.762
20	28.100
22	33.760
23	37.366
24	41.703
25	46.985
26	53.475
27	61.498
28	71.452
29	83.818
30	99.175
31	118.214
32	141.758
33	170.778
34	206.414
35	250.000

but there is variation in behaviour during the first part of the forecast window. There is evidence of initialisation shock in the SST field. The initial SST from the uncoupled ocean analysis is furthest from the true initial SST (~ 0.5 K warmer) and when the coupled model is initialised from the combined uncoupled atmosphere and ocean analysis states the forecast SST increases sharply, even further away from the true SST, over the first five model time-steps before gradually converging back towards the true trajectory. We also see jumps in the SST forecasts initialised from the strongly and weakly coupled analyses but these are much smaller suggesting that the coupled analyses are more balanced. Note that the differences between the SST forecasts can most easily be seen in Fig. 6, which focusses on the first 12 hours of the forecast window. In this example, the error in the weakly coupled SST analysis at the initial time is actually smaller than the strongly coupled SST analysis and the SST forecast from the weakly coupled analysis initially tracks the truth more closely. However, later in the forecast window, at the peak of the diurnal cycle (~ 25 hours), the SST forecasts from both the weakly and uncoupled analyses unexpectedly diverge from the truth, whereas the strongly coupled analysis continues to track it closely. This could be interpreted as a further indication of greater balance in the strongly coupled analysis; although the initial error in the SST forecast from the strongly coupled analysis is greater than the weakly

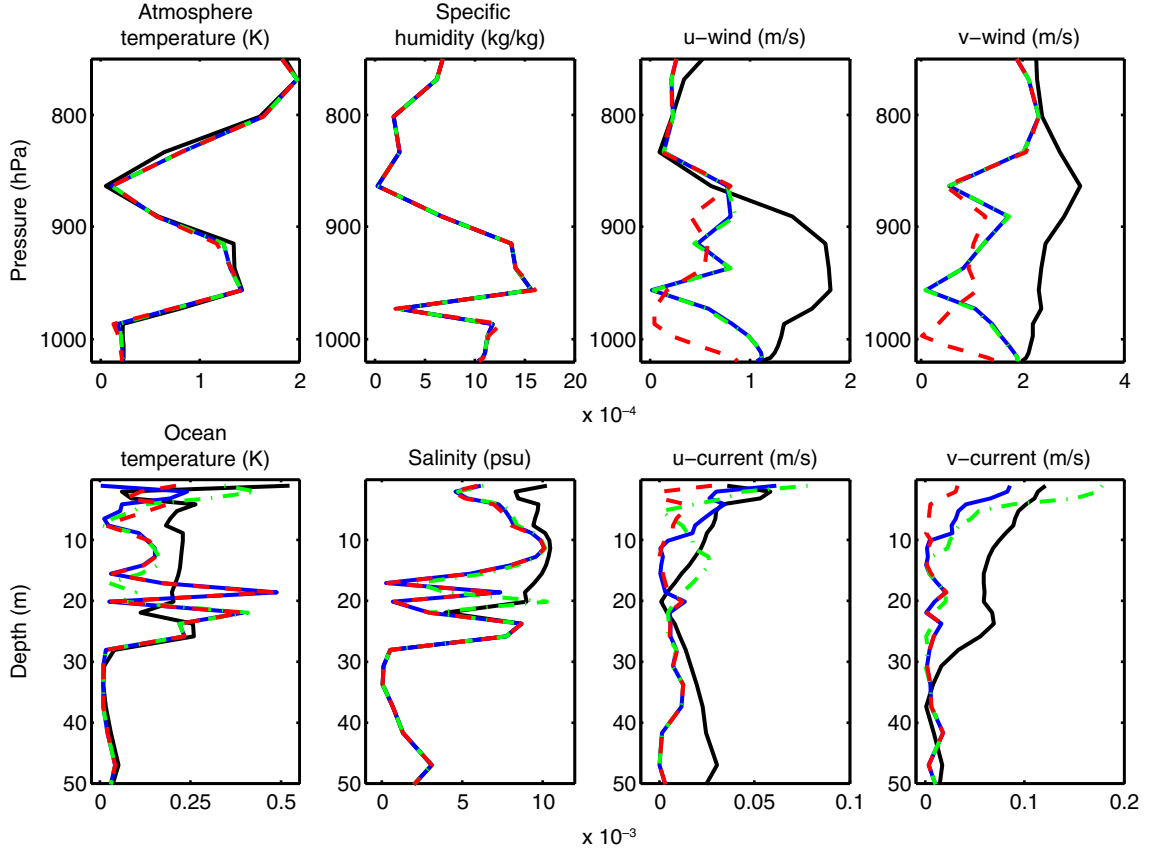


Fig. 3. Absolute errors at initial time: background error (solid black line), strongly coupled analysis error (dashed red line), weakly coupled analysis error (solid blue line), and uncoupled analysis error (dot-dash green line). Experiments using a 12-hour assimilation window with 3-hourly atmosphere & 6-hourly ocean observations. Observation locations are given in Tables 2 and 3.

coupled, for this example, it appears to be in better balance with rest of the model. The error in the initial temperature at the bottom atmosphere level is very similar for all three forecasts, but if we examine the atmosphere–ocean temperature difference (Fig. 5), we see that although at times the weakly and uncoupled systems are closer to the truth than the strongly coupled system, the temperature difference for the strongly coupled system is more stable across the whole forecast period. It also tracks the truth very accurately during the first 12 hours of the forecast which is the period corresponding to the assimilation window. The pattern seen in the sensible heat flux forecasts in Fig. 4 would also support this.

The strongly coupled analysis also produces better forecasts of the surface wind speed and u , v wind stress components (Fig. 4). The forecasts initialised from the weakly coupled and uncoupled analyses capture the general phasing of these fields but their magnitudes are overestimated to a greater extent than in the strongly coupled case over the first 24–48 hours of the forecast.

The latent heat flux forecasts are the slowest to stabilise; this is likely to be due to the fact that we are not assimilating

observations of specific humidity. However, as the forecasts adjust towards the model attractor, there is a clear pattern of increasing accuracy as we progress from uncoupled to weakly to strongly coupled initialisation.

Overall our experiments have shown that, when compared to uncoupled initialisation, initialisation using the analysis from a coupled assimilation can help to reduce initialisation shock and its impact on the subsequent forecast. For our model, the benefit appears to be greatest with the strongly coupled system; the weakly coupled assimilation system is also capable of reducing shock, but its behaviour is less consistent. It is worth noting that we would expect the type and size of shocks produced by each assimilation system to vary depending on factors such as season, location, assimilation start time and initial background state, and whether a full diurnal cycle is being observed.

5.2. Uncoupled assimilations with ‘true’ SST and surface fluxes

To understand how much the accuracy of the prescribed SST and surface fluxes may affect the results of the

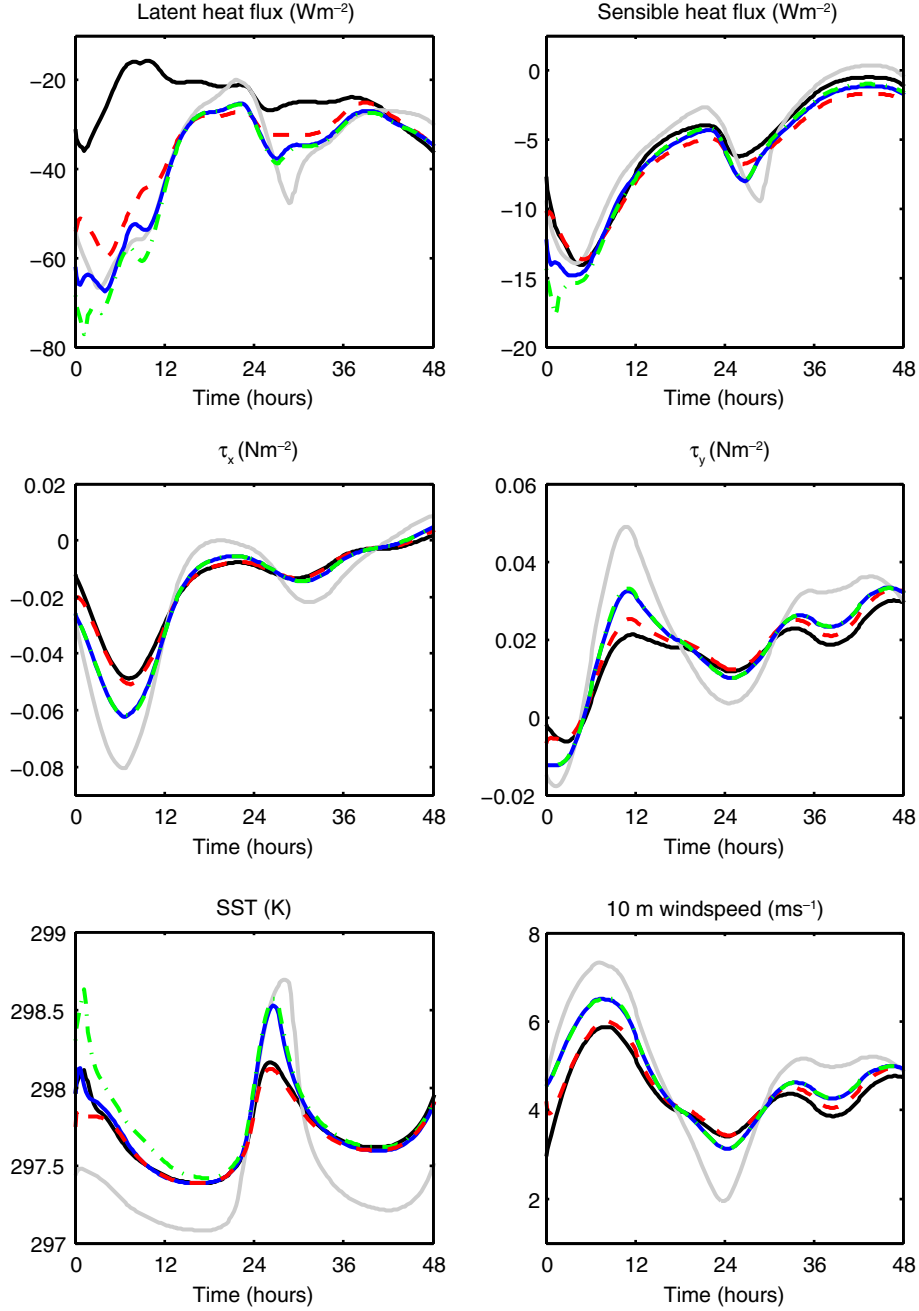


Fig. 4. Coupled model SST & surface fluxes for coupled model forecast initialised from t_0 analyses, first 48 hours of forecast: truth (solid black line), forecast initialised from initial background state (solid grey line), forecast initialised from strongly coupled analysis (dashed red line), forecast initialised from weakly coupled analysis (solid blue line), and forecast initialised from uncoupled analyses (dot-dash green line). Experiments using a 12-hour assimilation window with 3-hourly atmosphere & 6-hourly ocean observations.

uncoupled assimilation, we repeat the experiments performed in Section 5.1 using 6-hourly snapshots of the fluxes from the ‘truth’ trajectory in place of the ERA interim forcing fields. In some sense, this is the best we may expect the uncoupled assimilation system to be. Figure 6 shows the forecast fluxes for this case alongside those from the

original experiment for the first 12 hours of the forecast. Although we see an improvement, with reduced shocks in the SST and latent and sensible heat fluxes, the strongly coupled assimilation still generally performs better than the uncoupled assimilation. Even with the true forcing data, the uncoupled systems suffer from the lack of




Fig. 5. Atmosphere–Ocean temperature difference ($T_{60}-SST$) for coupled model forecast initialised from t_0 analyses: truth (solid black line), forecast initialised from strongly coupled analysis (dashed red line), forecast initialised from weakly coupled analysis (solid blue line), and forecast initialised from uncoupled analyses (dot–dash green line). Experiments using a 12-hour assimilation window with 3-hourly atmosphere & 6-hourly ocean observations.

atmosphere–ocean feedback. The weakly coupled SST analysis still produces a better SST forecast than the uncoupled ocean SST analysis, but the corresponding surface flux forecasts are either very similar or slightly worse for the weakly coupled case. This was also verified in experiments with different test cases (not shown). Overall, the performance of our weakly coupled assimilation system is usually comparable to the uncoupled system with the ‘true’ forcing. This indicates that even moving to a weakly coupled assimilation system may be of benefit. Furthermore, if multiple outer-loops are used, the update of the SST and surface fluxes can provide useful information not available to the uncoupled systems, as we demonstrate in Sections 5.4 and 5.5.

5.3. Temporal frequency of observations

To test the sensitivity of the different approaches to the frequency of observations, we repeat the assimilation experiments with the frequency of the atmosphere observations reduced by half to 6-hourly, so that we are observing the atmosphere and ocean with the same frequency. The strongly coupled system still performs better than the weakly coupled and uncoupled systems. The most significant effect is that differences between forecasts initialised from the weakly coupled and uncoupled analyses are less pronounced over the first 12 hours. This is best illustrated through the SST and surface flux forecasts (Figs. 7 and 8). If we compare these with Figs. 4 and 6, we see that both the strongly and weakly coupled t_0 SST estimates are further from the true value than in the previous case. The forecast initialised from the strongly coupled analysis adjusts itself smoothly, but the forecast initialised from the weakly coupled analysis exhibits a much larger shock

with amplitude similar to the uncoupled case. There is also more drift in the SST forecasts in the second half of the forecast window for this case.

The change in the SST trajectories means that there is now also less of a clear gap between the latent and sensible heat fluxes for the forecasts initialised from the uncoupled and weakly coupled analyses. There is no real change to the pattern of behaviour in the wind stresses, but the over estimation of magnitude is slightly larger due to the increased errors in the near-surface wind forecasts.

These experiments have shown that the weakly coupled assimilation system appears to be much more sensitive to the observation frequency than the strongly coupled system. This is because the weakly coupled assimilation system is, unlike the strongly coupled system, predominantly exposed to the coupling of the atmosphere and ocean through the innovations which are by definition in observation space. Therefore, if the number of observations is decreased, either spatially or temporally, this will clearly impact the greatest on the weakly coupled assimilation system. Although weak coupling can reduce shock, as seen in Section 5.1, at its worst it can produce results that are very similar to the uncoupled assimilation system. Similar behaviour was found in experiments varying the background error variances (not shown). Since changing the background error standard deviations has the effect of changing the relative weight given to the observations, the performance of the weakly coupled system was found to be more sensitive than the strongly coupled system to a change in the prescribed background error standard deviations.

5.4. Single observation experiments

A big hope for coupled data assimilation systems is that they will enable greater use of near-surface observations, such as satellite SST measurements and scatterometer data, by allowing cross-covariance information between the atmosphere and ocean. We investigate this property by assimilating single observations of near-surface variables. Since the initial background error covariance matrix is diagonal, increments from a single observation at the end of the assimilation window provide insight into the implicit covariances generated by the 4D-Var system. The purpose of the experiments presented here is to illustrate the ability of the strongly and weakly coupled systems to induce cross-correlations between the atmosphere and ocean rather than investigating their size and structure.

To understand the impact of SST observations, we assimilate the temperature from the top ocean level (1 m depth) at the end of the 12-hour assimilation window. Figure 9 shows the (analysis-background) increments produced by the strongly and weakly coupled systems at initial time, t_0 . We see initial increments in atmospheric

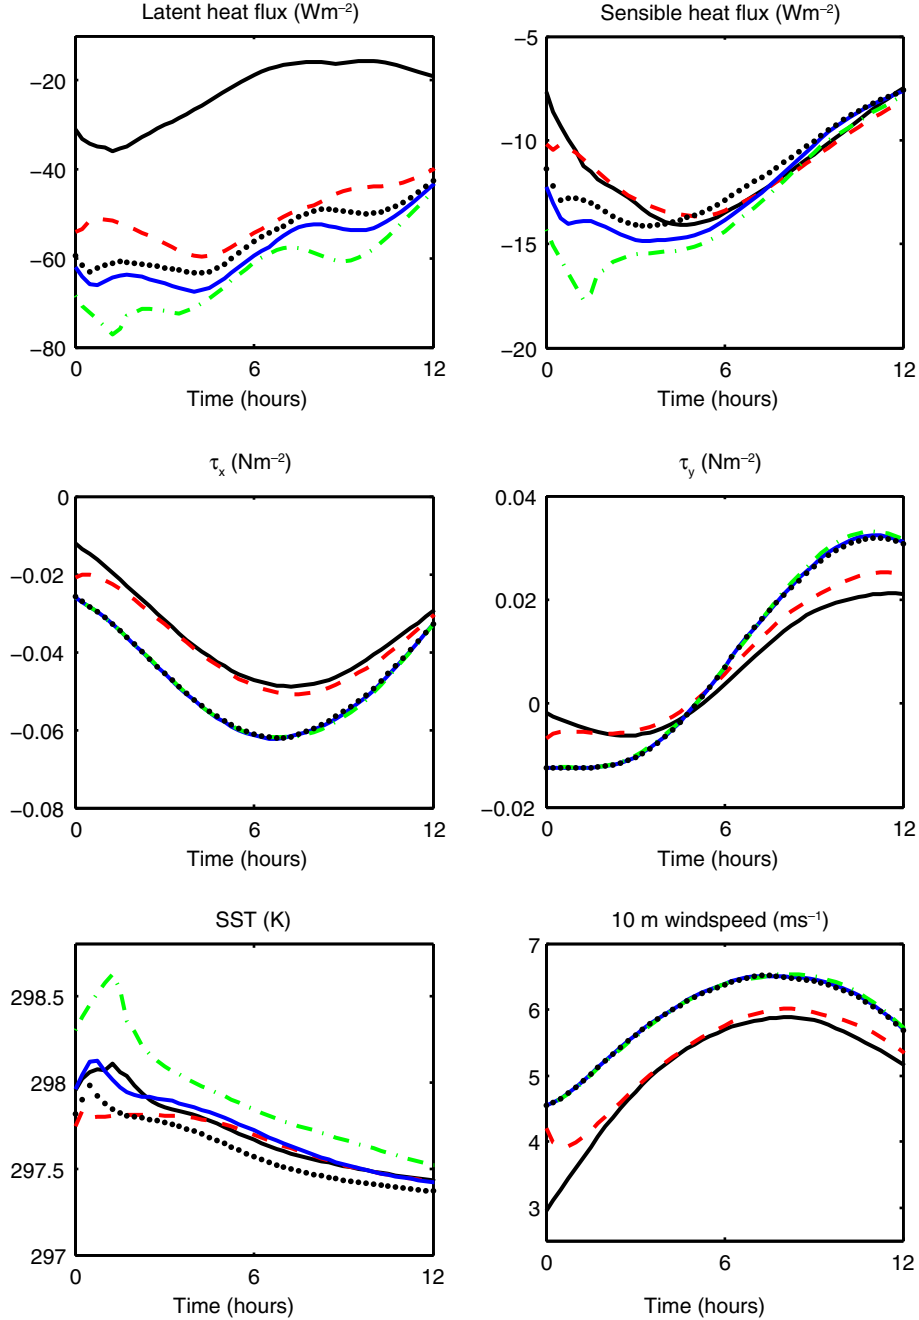


Fig. 6. Coupled model SST & surface fluxes for coupled model forecast initialised from t_0 analyses, first 12 hours of forecast: truth (solid black line), forecast initialised from strongly coupled analysis (dashed red line), forecast initialised from weakly coupled analysis (solid blue line), forecast initialised from uncoupled analyses (dot-dash green line), and forecast initialised from uncoupled analyses using ‘true’ forcing (black dots). Experiments using a 12-hour assimilation window with 3-hourly atmosphere & 6-hourly ocean observations.

temperature, specific humidity, ocean temperature and salinity with the strongly coupled system, but only ocean temperature and salinity for the weakly coupled system.

We can relate the behaviour observed in this experiment back to the model equations and assimilation system design. The strongly coupled system uses the coupled TL

and adjoint models in its inner-loops. Since the boundary conditions for the atmospheric temperature and specific humidity depend on the SST and the boundary conditions for the ocean temperature and salinity depend on the atmospheric temperature and humidity via the latent and sensible heat fluxes, an increment or perturbation to the

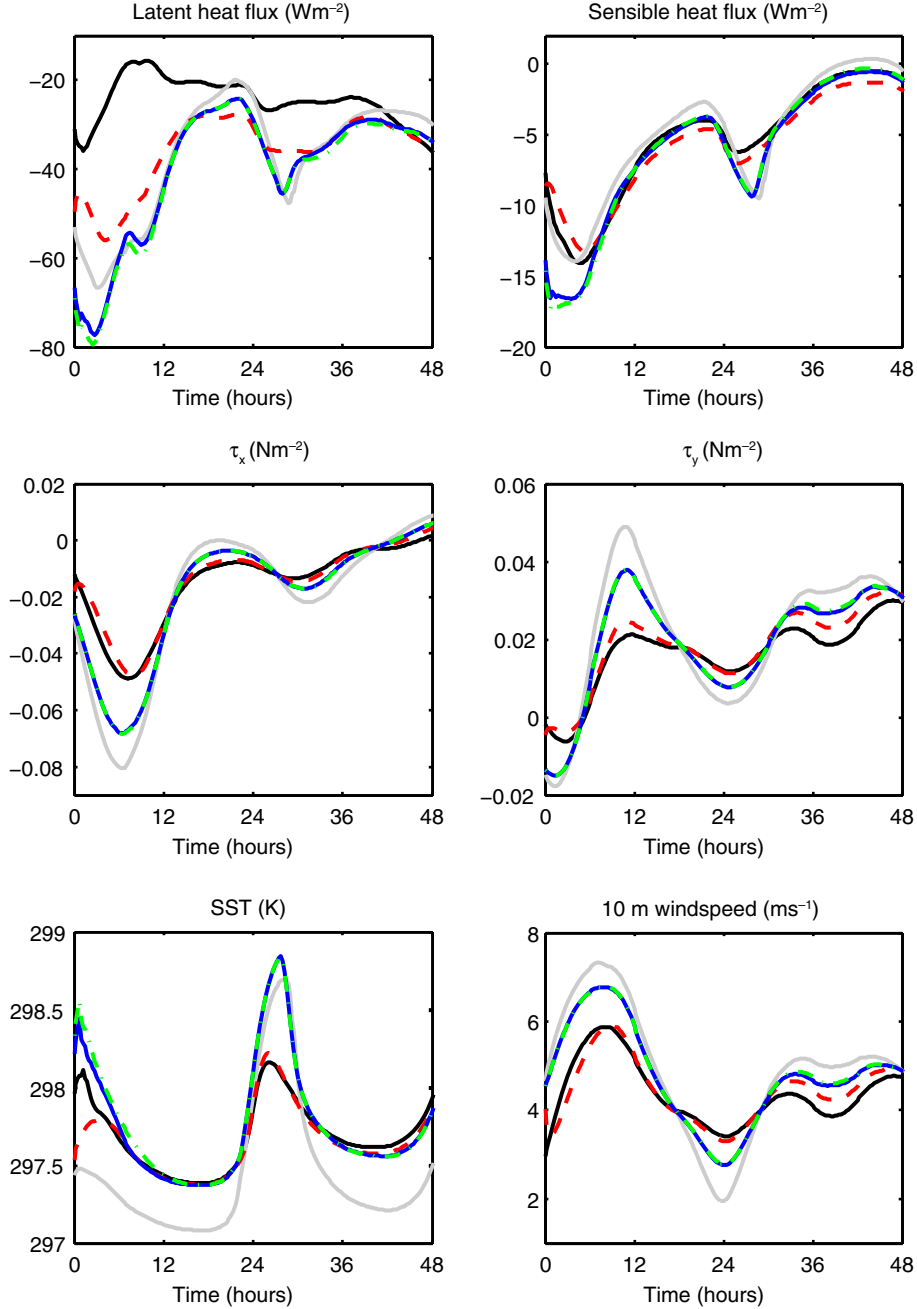


Fig. 7. Coupled model SST & surface fluxes for coupled model forecast initialised from t_0 analyses, first 48 hours of forecast: truth (solid black line), forecast initialised from initial background state (solid grey line), forecast initialised from strongly coupled analysis (dashed red line), forecast initialised from weakly coupled analysis (solid blue line), and forecast initialised from uncoupled analyses (dot-dash green line). Experiments using a 12-hour assimilation window with 6-hourly atmosphere & 6-hourly ocean observations.

SST should produce increments in the atmosphere temperature and specific humidity fields. We do not see increments in the initial u and v wind fields because assumptions made in the development of the TL and adjoint models mean that the SST does not directly depend on them. The weakly coupled system cannot produce initial

increments in the atmosphere fields because it runs separate inner-loops for the atmosphere and ocean which use the uncoupled atmosphere-only and uncoupled ocean-only TL and adjoint models. The SST used in the boundary conditions for the inner-loop uncoupled atmosphere model and the surface fluxes used in the boundary conditions of

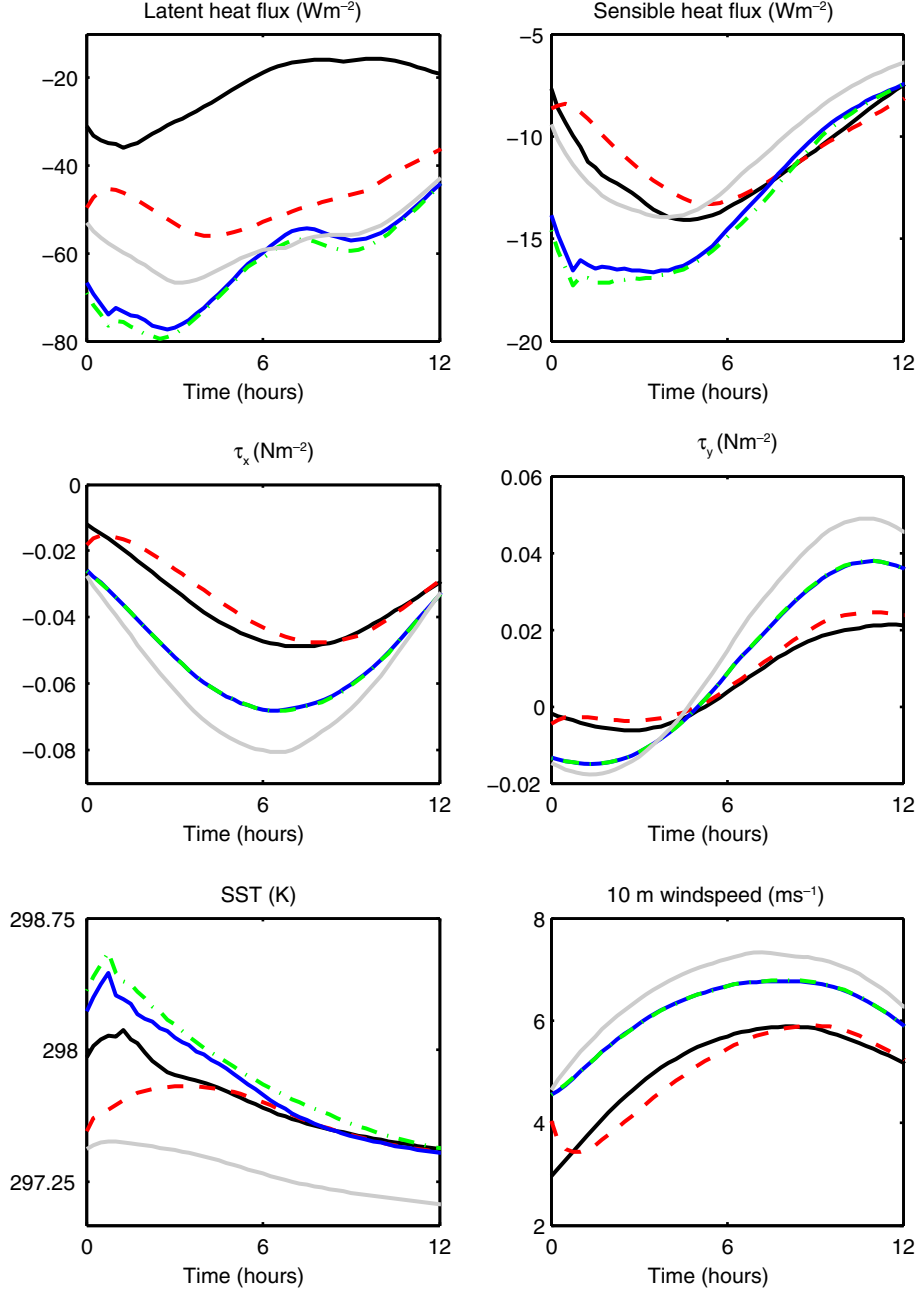


Fig. 8. Coupled model SST & surface fluxes for coupled model forecast initialised from t_0 analyses, first 12 hours of forecast: truth (solid black line), forecast initialised from initial background state (solid grey line), forecast initialised from strongly coupled analysis (dashed red line), forecast initialised from weakly coupled analysis (solid blue line), and forecast initialised from uncoupled analyses (dot-dash green line). Experiments using a 12-hour assimilation window with 6-hourly atmosphere & 6-hourly ocean observations.

the inner-loop uncoupled ocean model are prescribed from the outer-loop linearisation trajectory.

The initial analysis increments modify the coupled model trajectory and produce increments to the background fields for all variables across the rest of the assimilation window; Fig. 10 shows the analysis increments in the centre of the assimilation window ($t = 6$ hours) as an example. The changes in

the initial atmospheric temperature, specific humidity, ocean temperature and salinity fields subsequently produce increments to all the model variables. However, these increments are relatively small compared to the full fields and so when we examine the analysis trajectories across the whole assimilation window we see that the ocean u_o , v_o current and atmosphere u , v wind trajectories are qualitatively

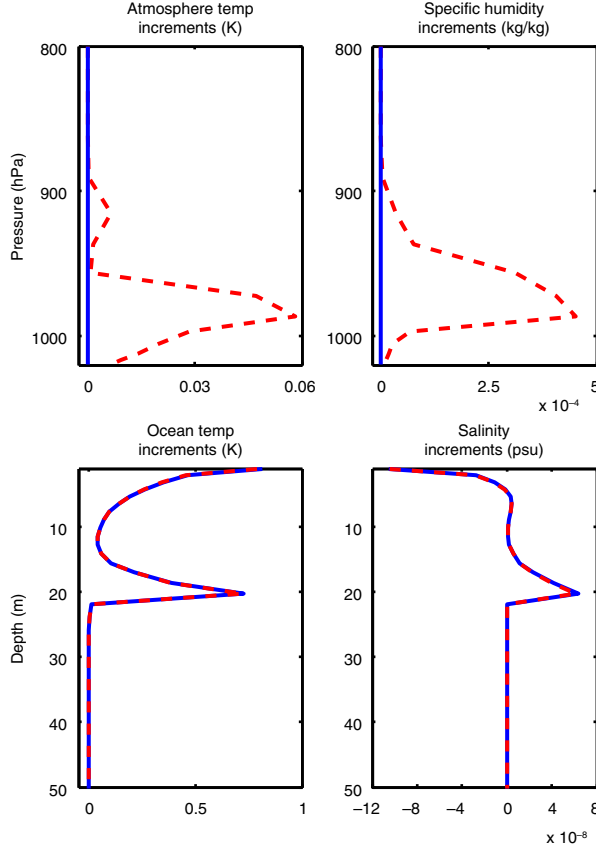


Fig. 9. Analysis increments at $t=0$: strongly coupled (dashed red line) and weakly coupled (solid blue line) assimilations with single SST observation at the end of a 12-hour assimilation window.

very similar for both the strongly and weakly coupled systems (not shown). There are more visible differences between the strongly and weakly coupled atmospheric temperature and specific humidity analysis trajectories due to the difference in increments at t_0 (not shown).

Scatterometer data provide information on ocean surface wind speed and direction via measurements of backscatter from surface waves. Since our system is only currently designed to handle direct observations we use (1) horizontal wind components, u and v , at the bottom level of the atmosphere model (~ 10 m height); (2) zonal and meridional ocean currents at the top level of the ocean model as a proxy.

With single u and v wind observations only the u and v wind fields are updated at t_0 for both the strongly and weakly coupled systems (not shown). These initial wind increments do, however, produce increments to all of the atmosphere and ocean background fields over the remainder of the assimilation window. In this case, the strongly and weakly coupled analyses are identical; this is due to the model formulation and the fact that we are ignoring

perturbations to the diffusion coefficients in the TL and adjoint models as described in Section 3.2. The u and v winds only depend on each other and so are essentially decoupled from the rest of the model in both the coupled model and atmosphere-only model.

With u_o and v_o ocean surface current observations, the strongly coupled system produces initial increments in all fields, although these are very small for atmospheric temperature and specific humidity. The weakly coupled system only produces initial increments in the ocean fields and these are larger than in the strongly coupled case (Fig. 11). Again, the update of the initial state gives rise to increments in all fields across the assimilation window for both systems. Although the t_0 ocean analysis increments are larger in the weakly coupled system, the increments across the assimilation window are generally smaller, particularly for the atmospheric fields, where there is no change to the initial states (results not shown). There is no difference in the strongly and weakly coupled SST and surface fluxes when we assimilate single wind observations and only very small differences in the single SST observation experiment. However, for this case the strongly coupled system produces a much better analysis of the true surface wind stress and wind speed than the weakly coupled system (Fig. 12). As described in Section 2.1.1, the strongly coupled system is able to generate cross-covariances between the atmosphere and ocean fields and thus improve the wind analysis using the ocean current observations. Improved near-surface wind conditions can have a positive impact on air–sea exchange and thus both the atmosphere and ocean analyses. This result clearly demonstrates the potential for greater use of near-surface data with strongly coupled assimilation.

These experiments have provided a valuable illustration of the ability of a strongly coupled assimilation system to induce cross-covariance information between the atmosphere and ocean variables, such that a single observation of a variable in one fluid at the end of the assimilation window can produce increments to variables in the other fluid at initial time t_0 . Although the structure of the weakly coupled assimilation system does not allow atmosphere–ocean cross-covariances, there is benefit to be gained from this approach if more than one outer-loop is used, and particularly if both the atmosphere and ocean are well observed (see Section 5.3). An analysis increment from an observation in one system will change the linearisation state for both the uncoupled TL models used in the next inner-loop minimisation and thus has the potential to influence the subsequent analysis across the whole atmosphere–ocean system. In the next section, we present a simple illustration of how assimilating a single observation of both fluids enables the weakly coupled system to modify the initial analysis increments across the air–sea interface.

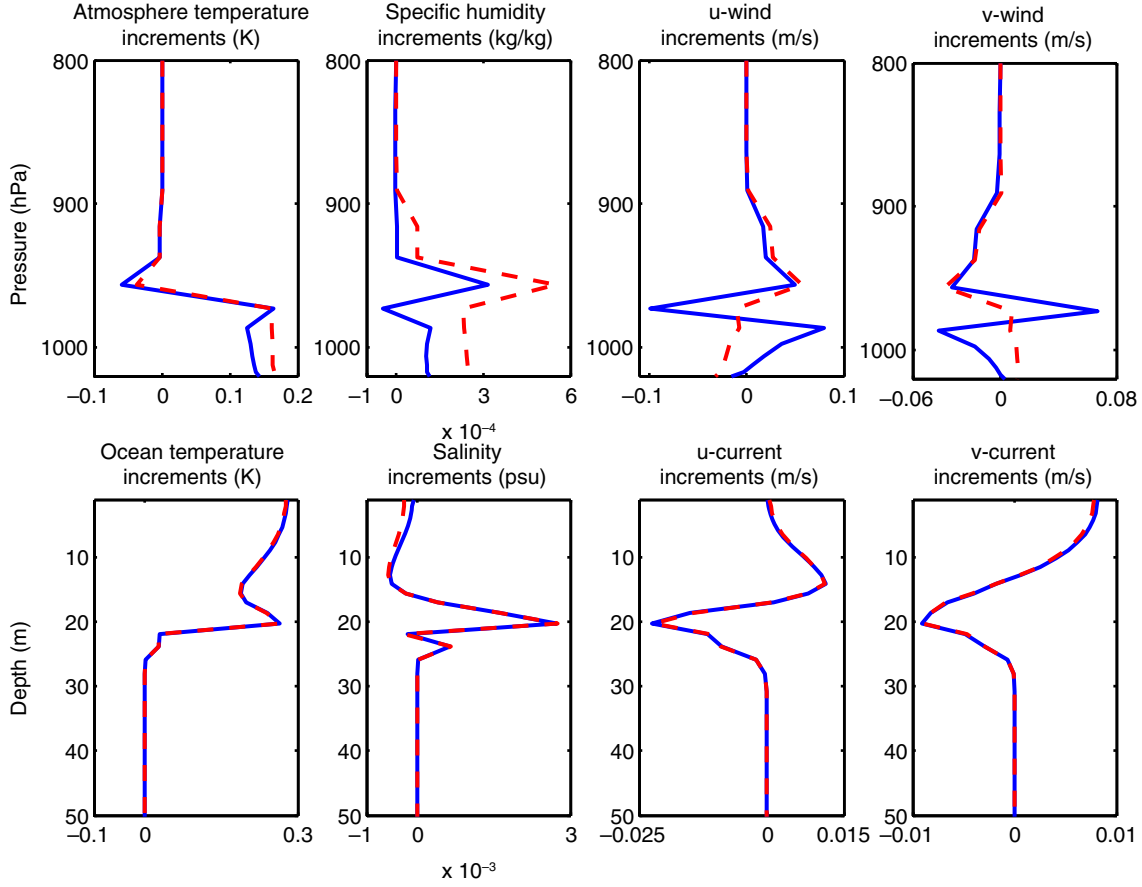


Fig. 10. Analysis increments at $t=6$ hours: strongly coupled (dashed red line) and weakly coupled (solid blue line) assimilations with single SST observation at the end of a 12-hour assimilation window.

5.5. Double observation experiment

Due to the inability of the weakly coupled system to produce an initial increment in the unobserved fluid it is not possible to see explicitly how information is spread across the air–sea interface in the analysis increments when only one fluid is observed. Unlike in the uncoupled system, in the weakly coupled system observations are able to influence the analysis across the atmosphere–ocean interface. This can be illustrated by assimilating two observations of temperature at the end of the 12-hour assimilation window; one at the lowest level of the atmosphere, and one at the top level of the ocean. The atmospheric temperature analysis increments after three outer-loops are illustrated in Fig. 13 (solid lines). These can be compared to the analysis increments when only the atmosphere temperature observation is assimilated (dashed lines) and when only the ocean temperature observation is assimilated (dot–dash lines, same as Fig. 9).

It can be seen in the strongly coupled case that only assimilating the atmosphere temperature observation produces a negative increment, peaking at approximately

980 hPa, and assimilating only the ocean temperature observation produces a positive increment, peaking again at approximately 980 hPa. When both these observations are assimilated, the strongly coupled system is able to make use of the overlapping information they provide and the result is a positive increment which is slightly reduced in magnitude compared to when only the ocean observation is assimilated. For the weakly coupled case, when only the atmosphere temperature observation is assimilated the atmospheric temperature analysis increment is nearly identical to that produced by the strongly coupled system. When only the ocean temperature observation is assimilated, there is no initial analysis increment in the atmosphere temperature field (as discussed in Section 5.4). However, when both temperature observations are assimilated the atmospheric temperature analysis increment becomes larger in magnitude than when only the atmosphere temperature observation was assimilated. Therefore, the weakly coupled system is able to use the ocean observation to alter the initial analysis increment in the atmosphere when atmosphere observations are also assimilated. This is not possible in the uncoupled system.

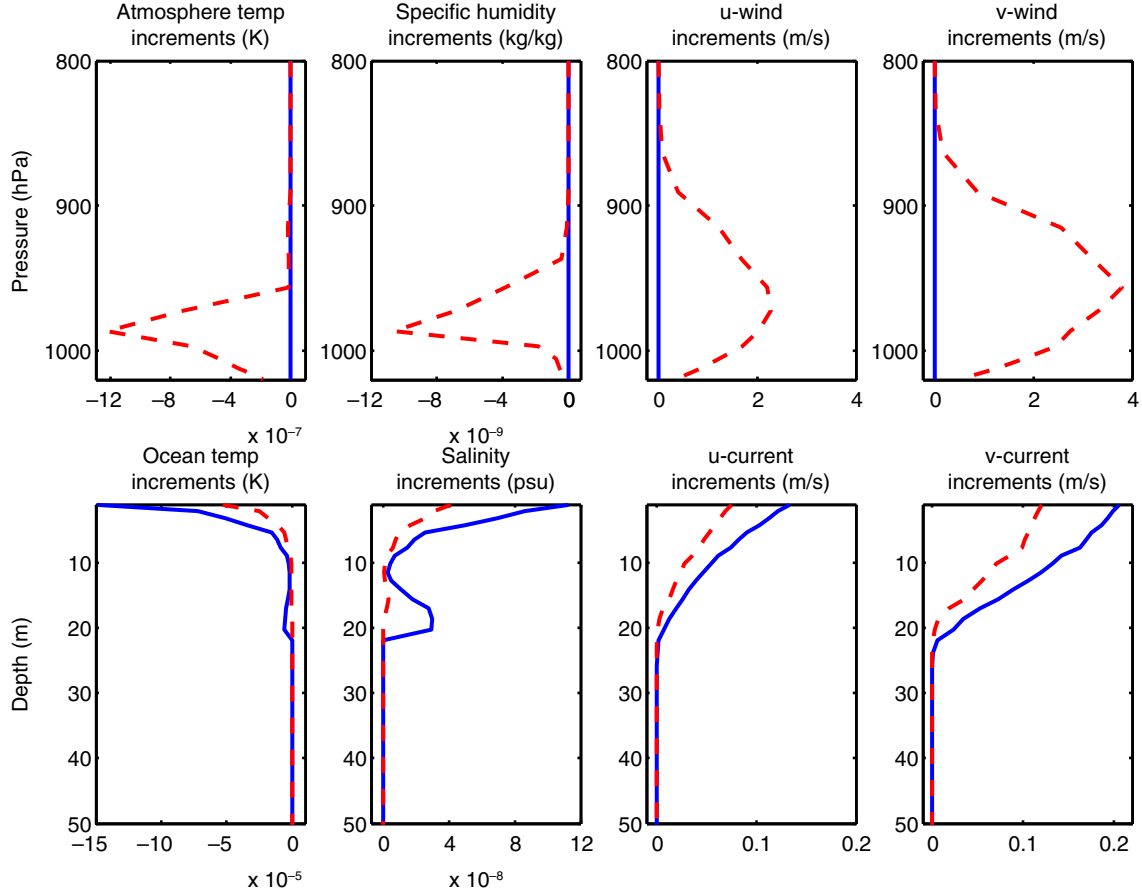


Fig. 11. Analysis increments at $t=0$: strongly coupled (dashed red line) and weakly coupled (solid blue line) assimilations with single ocean surface current observation at end of 12-hour assimilation window.

6. Summary

We have developed an idealised coupled atmosphere–ocean model system and used it to study different formulations of the coupled atmosphere–ocean data assimilation problem. By employing the incremental 4D-Var algorithm, we have built the capability to run both strongly and weakly coupled assimilations as well as uncoupled atmosphere- or ocean-only assimilations. This has provided a flexible framework for comparing the behaviours of varying degrees of coupling.

A key motivation for the development of coupled data assimilation systems is the potential for the reduction or elimination of initialisation shock in coupled model forecasts via the generation of more balanced initial conditions, and the positive impact this is expected to have in terms of forecast skill. Initialisation shocks were seen in SST in our simple system and experiments showed that, when compared to uncoupled initialisation, coupled assimilation is able to reduce initialisation shock and its impact on the subsequent forecast, although it may not eliminate it

completely. Whilst this improvement was clearly evident when using analyses from the strongly coupled system, it was not always so obvious with the weakly coupled system in our experiments. The ability of the weakly coupled assimilation system to reduce initialisation shock was found to be sensitive to the input parameters, such as observation frequency. In the best cases, the behaviour of the SST and surface fluxes in the initial stages of the forecast (used to identify shock) followed those from the strongly coupled assimilation. In other cases, the weakly coupled assimilation did not show the same improvement as strongly coupled assimilation. However, the weakly coupled system was usually comparable to uncoupled assimilations in which the atmosphere and ocean models were forced using the ‘true’ SST and surface fluxes. This illustrates that even moving to a weakly coupled assimilation system should be of benefit, as the update of the SST and surface fluxes through the outer-loop step can provide useful information not available to the uncoupled assimilation systems.

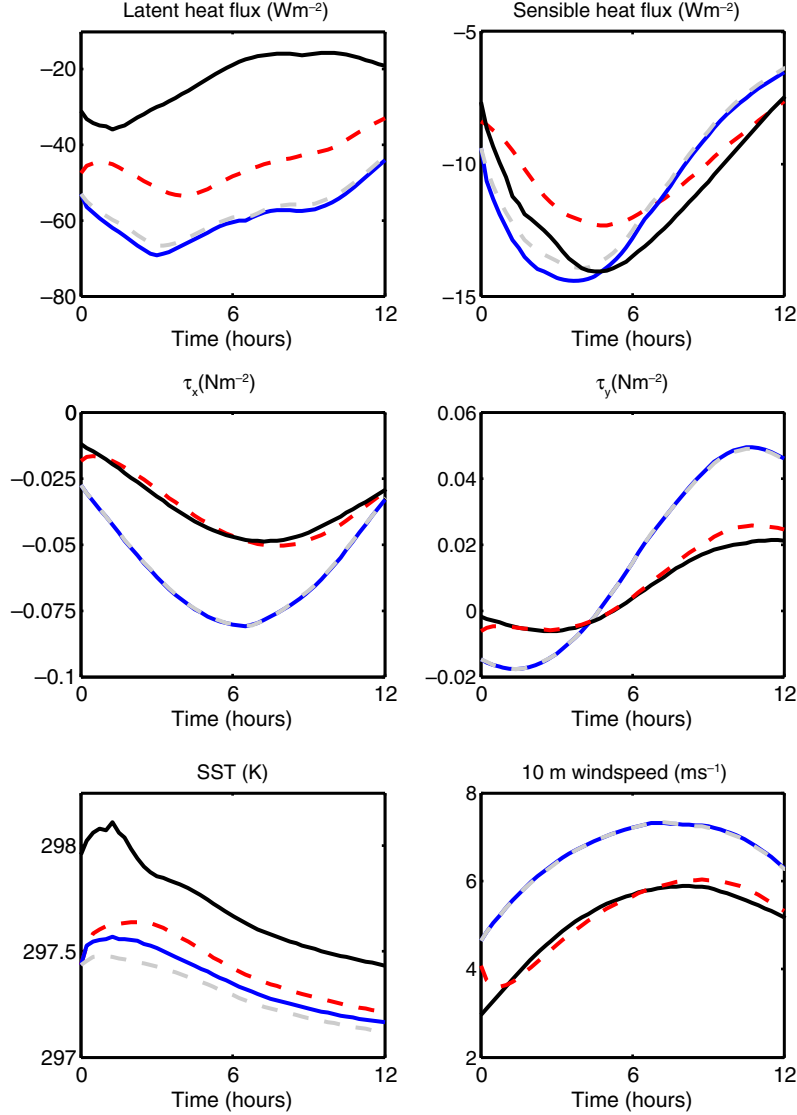


Fig. 12. SST & surface fluxes: truth (solid black line), initial background trajectory (dashed grey line), strongly coupled analysis (dashed red line) and weakly coupled analysis (solid blue line) from experiment with single ocean surface current observation at the end of a 12-hour assimilation window.

Single observation experiments were used to demonstrate how coupled assimilation systems offer the potential for improved use of near-surface observations via the generation of cross-covariance information. Although the possible cross-covariances that can be generated are partly limited by the simplified dynamics of our model, the effect of coupled assimilation can clearly be seen. The strongly coupled assimilation system is able to implicitly induce cross-covariance information between the atmosphere and ocean at the initial time, such that a single ocean observation can generate analysis increments in the initial atmospheric fields and vice versa. While the design of the weakly coupled incremental 4D-Var assimilation algorithm does

not allow this, the use of the coupled model in the outer-loop update step means that if more than one outer-loop is run, an observation in one system can affect the other system by changing the linearisation state. With observations in both fluids, the weakly coupled system is also able to alter the analysis increments across the atmosphere–ocean interface at the initial time and this was illustrated via a double observation experiment. Thus information from near-surface observations can be used to greater effect compared to uncoupled assimilation systems.

Overall, the results from experiments with this idealised system support the belief that benefits can be expected from coupled data assimilation systems. In the experiments




Fig. 13. Atmospheric temperature analysis increments at $t=0$: strongly coupled (left) and weakly coupled (right) assimilations with two observations of temperature, corresponding to the lowest model level in the atmosphere and the top model level in the ocean, at the end of a 12-hour assimilation window. The solid lines indicate the analysis increment when both temperature observations are assimilated; the dashed lines correspond to assimilation of the atmosphere temperature observation only; and dot-dash lines correspond to assimilation of the ocean temperature observation only.

presented here and others performed using variations of the set-up described, the strongly coupled assimilation system generally outperforms both the weakly coupled and uncoupled systems, in terms of producing more balanced initial analysis fields, and extracting more information from observations through the implicit generation of cross-covariances. The results from the weakly coupled assimilation experiments show that benefit can be gained from such a system, but that it is unlikely to be as large as that from strongly coupled assimilation. Nevertheless, even with a weak coupling we may expect some reduction in initialisation shock and the generation of some cross-covariance information. Thus, the current efforts of operational centres to develop weakly coupled assimilation systems are a step in the right direction.

Further work is required to better understand the sensitivity of the weakly coupled system to the input parameters of the assimilation. In particular, this study used a diagonal background error covariance matrix in order to understand more cleanly the covariances generated by the coupled assimilation. If the weakly coupled assimilation included non-diagonal background error covariance matrices in the atmosphere and ocean inner-loop cost functions, then better balance would be expected in the increments of the individual systems. This may in itself

help to reduce initialisation shock and make better use of observations.

Work is now underway to investigate the nature and structure of the atmosphere–ocean cross-covariances and how they should be represented in both strongly and weakly coupled systems. An increased understanding of the covariance information arising from atmosphere–ocean coupling will provide valuable guidance for the design of more balanced covariances for future full-scale coupled data assimilation systems.

7. Acknowledgements

This work was funded by the European Space Agency (ESA) as part of the Data Assimilation projects – Coupled Model Data Assimilation initiative (contract no. 4000104980/11/I-LG), and the UK Natural Environment Research Council (NERC) via research grant no. NE/J005835/1 and as part of NERC’s support of the National Centre for Earth Observation (NCEO). We are also grateful to Keith Haines and colleagues at the ECMWF and UK Met Office for useful discussions, and to Irina Sandu and Glenn Carver of the ECMWF for providing source code and technical support. We would like to thank the anonymous reviewers for their helpful comments and in particular the reviewer who suggested the experiment presented in Section 5.5.

8. Appendix

Atmosphere and ocean model components

This appendix provides further details of the atmosphere and ocean components of the coupled 1D model system.

A.1. Atmosphere model

The atmospheric component of the model is a stripped-down version of the ECMWF single-column model which originates from an early cycle of the IFS (Integrated Forecasting System) code. The model solves the primitive equations for temperature, T , specific humidity, q , and zonal, u , and meridional, v , wind components, formulated in non-spherical co-ordinates, (Simmons and Burridge, 1981; Ritchie et al., 1995) and using a hybrid vertical co-ordinate, η (see Section A.1.2),

$$\frac{\partial u}{\partial t} + \dot{\eta} \frac{\partial u}{\partial \eta} = f(v - v_g) + F_u + P_u, \quad (\text{A1})$$

$$\frac{\partial v}{\partial t} + \dot{\eta} \frac{\partial v}{\partial \eta} = -f(u - u_g) + F_v + P_v, \quad (\text{A2})$$

$$\frac{\partial T}{\partial t} + \dot{\eta} \frac{\partial T}{\partial \eta} = \frac{R}{c_p} T \frac{\omega}{p} + F_T + P_T, \quad (\text{A3})$$

$$\frac{\partial q}{\partial t} + \dot{\eta} \frac{\partial q}{\partial \eta} = F_q + P_q. \quad (\text{A4})$$

Here t is time, f is the Coriolis parameter, u_g and v_g are prescribed geostrophic wind components, p is pressure, R is the gas constant for air, c_p is the specific heat at constant pressure for air, $\dot{\eta}$ is the vertical velocity in η co-ordinates and ω is the prescribed vertical velocity in pressure co-ordinates. The F_ϕ ($\phi = u, v, T, q$) are forcing terms representing the horizontal advection of the mean variables and the P_ϕ terms represent tendencies due to the parameterisation of sub-grid scale physical processes.

The vertical advection terms in eqs. (A1)–(A4) are computed using a two time level Eulerian (upwind) scheme with a semi-implicit treatment of the right hand sides (ECMWF IFS documentation, 2001–2013a).

In the original ECMWF SCM code (the ‘full physics’ version), the P_ϕ terms in eqs. (A1)–(A4) incorporate the effects of processes such as radiation, turbulent mixing, moist convection and clouds. For the simplified system, the code was stripped back to include just advection and turbulent mixing; the P_ϕ terms then represent physical tendencies due to vertical exchange by turbulent processes only. This was done in order to simplify the derivation of the adjoint model whilst ensuring that the evolution of the atmosphere was sufficiently realistic for purposes of this study. The turbulent mixing is parameterised using a k -diffusion approach (Louis, 1979)

$$\frac{\partial \phi}{\partial t} = \frac{1}{\rho_a} \frac{\partial J_\phi}{\partial z}, \quad (\text{A5})$$

where ϕ is the prognostic variable (T, q, u , or v), ρ_a is the density of air, z is height, and the vertical turbulent flux J_ϕ (positive downward) is given by

$$J_\phi = \rho_a K_\phi \frac{\partial \phi}{\partial z}, \quad (\text{A6})$$

where K_ϕ is the turbulent exchange coefficient.

The exchange coefficients between the surface and lowest model level (~ 10 m above surface) are expressed as functions of the bulk Richardson number, determined according to the formulation of Louis et al. (1982). Above the surface layer, the turbulent transports are based on local stability and the coefficients are defined using a combined Louis–Tiedtke–Geleyn (LTG)–Monin–Obukov (MO) formulation (Beljaars, 1995; Beljaars and Viterbo, 1998; Viterbo et al., 1999). The physical tendencies are computed using an implicit time-stepping procedure. Full details of the turbulent diffusion scheme together with further references can be found in the ECMWF IFS documentation (2001–2013b).

A.1.1. Boundary conditions

The upper and lower boundary conditions are given by

$$K_\phi \frac{\partial \phi}{\partial z} = 0, \quad \text{at } p = p_{\text{top}} \quad (\text{A7})$$

$$K_\phi \frac{\partial \phi}{\partial z} \rightarrow C_\phi |U(z)| (\phi(z) - \phi_{\text{surf}}) \quad \text{as } z \rightarrow 0. \quad (\text{A8})$$

where p_{top} is the pressure at the top of the atmosphere (set at 0.1 hPa), C_ϕ is the transfer coefficient at the lowest model level, and ϕ_{surf} represents the value of the prognostic variable ϕ at the surface. The SST and saturation specific humidity are used as ϕ_{surf} values for temperature and specific humidity, and a no-slip condition is used for the u and v wind components. The surface turbulent fluxes are passed to the ocean model where they are used in the computation of the ocean surface boundary conditions (Section A.2.2).

A.1.2. Vertical discretisation

The atmosphere is divided into 60 unequally spaced layers, extending from the surface up to p_{top} , with the finest vertical resolution (measured in geometric height) in the planetary boundary layer. The model uses the hybrid vertical co-ordinate of Simmons and Burridge (1981). The co-ordinate $\eta = \eta(p, p_s)$ is a monotonic function of the pressure, and is also dependent on the surface pressure, p_s . There is no staggering of prognostic model variables; T, q, u and v are all represented at ‘full-level’ pressures p_k and the model layers are defined by the pressures at the interfaces between them (termed ‘half levels’).

A.2. Ocean mixed layer model

The ocean mixed layer model is based on the KPP vertical mixing scheme of Large et al. (1994). The code was originally developed by the NCAS Centre for Global Atmospheric Modelling at the University of Reading (Woolnough et al., 2007) and incorporated into the ECMWF SCM code by Takaya et al. (2010) as part of a study into the impact of better representation of coupled atmosphere–upper ocean processes in the ECMWF medium-range forecasts. In this section, we summarise the components of the scheme most relevant to this study; a comprehensive description of the model is given in Large et al. (1994).

The KPP model describes the evolution of the mean values of temperature, θ , salinity, s , and zonal and meridional currents u_o, v_o . The time evolution of each field is expressed as the vertical divergence of the kinematic turbulent fluxes, $\overline{w'\phi'}$ ($\phi = \theta, s, u_o, v_o$), giving the following set of equations

$$\frac{\partial \bar{\phi}}{\partial t} = - \frac{\partial \overline{w'\phi'}}{\partial z} - \frac{\partial Q_n}{\partial z}, \quad (\text{A9})$$

$$\frac{\partial \bar{s}}{\partial t} = -\frac{\partial \overline{w's'}}{\partial z}, \quad (\text{A10})$$

$$\frac{\partial \bar{u}_o}{\partial t} = -\frac{\partial \overline{w'u'_o}}{\partial z} + f v_o, \quad (\text{A11})$$

$$\frac{\partial \bar{v}_o}{\partial t} = -\frac{\partial \overline{w'v'_o}}{\partial z} - f u_o. \quad (\text{A12})$$

Here, an overbar denotes a time average, primed variables represent turbulent fluctuations from this average, w is the turbulent vertical velocity and Q_n is the non-turbulent heat flux (solar irradiance) which is modelled using an empirical function of short wave radiation, Q_{SW} , and ocean depth, d (distance from ocean surface boundary).

The ocean surface boundary layer is defined as the region where d is less than or equal to the ocean boundary layer depth h , the value of which is based on the depth at which the bulk Richardson number equals the prescribed critical Richardson number. Within this region, the kinematic fluxes $\overline{w'\phi'}$ are parameterised using K-profiles

$$\overline{w'\phi'} = -K_\phi \frac{\partial \bar{\phi}}{\partial z}, \quad (\text{A13})$$

where $\bar{\phi}$ represents a mean quantity. The K_ϕ are expressed as the product of a depth dependent turbulent velocity scale and a smooth non-dimensional shape function such that they are directly proportional to h at all depths.

In the ocean interior ($d > h$) the turbulent vertical fluxes are parameterised as

$$\overline{w'\phi'} = -\nu_\phi(d) \frac{\partial \bar{\phi}}{\partial z}, \quad (\text{A14})$$

where the interior diffusivity ν_ϕ is the sum of resolved shear instability and unresolved shear instability due to internal wave breaking; we neglect the effect of double diffusion for reasons described in Takaya et al. (2010).

A.2.1. Geostrophic currents. For a 1D water column, the ocean currents are essentially governed by Ekman flow (Stewart, 2008). Without pressure gradient terms, the water moves under the sole influence of the Coriolis force and the ocean momentum equations reduce to the equation for the harmonic oscillator (Stewart, 2008); the solution takes the form of an inertial oscillation or inertial current. In reality, we expect the ocean currents to be approximately geostrophically balanced. To alleviate the unrealistic behaviour that this produces, we use the method of Takaya et al. (2010) and decompose the currents into slow and fast varying flows. The fast varying flow is assumed to be mainly the Ekman (or ageostrophic) flow simulated by the KPP model. The slow varying geostrophic component is prescribed and not modelled.

A.2.2. Boundary conditions. The ocean surface boundary conditions are given by the surface kinematic fluxes of heat, salt and momentum

$$\overline{w'\theta'_0} = -Q_t/(\rho_0 c_{p0}), \quad (\text{A15})$$

$$\overline{w's'_0} = -F_t s_0/\rho_0(0), \quad (\text{A16})$$

$$\overline{w'u'_0} = -\tau_x/\rho_0, \quad (\text{A17})$$

$$\overline{w'v'_0} = -\tau_y/\rho_0, \quad (\text{A18})$$

where Q_t is the net turbulent heat flux, F_t is the net turbulent freshwater flux, τ_x and τ_y are the zonal and meridional components of the surface wind stress, s_0 , ρ_0 , c_{p0} are the salinity, density and specific heat at constant pressure at the ocean surface, and $\rho_0(0)$ is the density of surface water with zero salinity (i.e. pure water). The fluxes Q_t and F_t are computed as

$$Q_t = Q_{LW} + Q_E + Q_H, \quad (\text{A19})$$

$$F_t = Q_E/L_v, \quad (\text{A20})$$

where Q_{LW} is net long wave radiation, Q_E , Q_H are the latent and sensible heat fluxes and L_v is the latent heat of evaporation. The latent and sensible heat and momentum fluxes are the surface turbulent fluxes from eq. (A8); these are computed within the atmosphere model component using the formulae given in Section A.3.

A.2.3. Vertical discretisation. The ocean model uses a stretched vertical grid (Takaya et al., 2010) with 35 levels from the surface to a depth of 250 m. The resolution is increased in the upper layers in order to simulate the diurnal SST variability; the top model layer is chosen to be 1 m thick and there are 19 levels in the top 25 m. The largest depth is fixed so that the ocean model levels do not vary with time. As with the atmosphere component, there is no staggering of the prognostic model variables; θ , s , u_o and v_o are all represented at full model level depths.

A.3. Atmosphere–ocean coupling

The atmosphere–ocean fluxes are estimated from bulk formulae

$$\tau_x = \rho_a C_D |U_n| u_n, \quad (\text{A21})$$

$$\tau_y = \rho_a C_D |U_n| v_n, \quad (\text{A22})$$

$$Q_H = \rho_a C_H |U_n| (T_n - SST), \quad (\text{A23})$$

$$Q_E = \rho_a L_v C_E |U_n| (q_n - q_{sat}(SST)), \quad (\text{A24})$$

where the subscript n represents the lowest atmosphere model level,

$$|U_n| = \sqrt{u_n^2 + v_n^2}, \quad (\text{A25})$$

is the (~ 10 m) windspeed and $q_{sat}(SST)$ is the surface saturation specific humidity. The drag coefficient, C_D , and the transfer coefficients for heat, C_H , and moisture, C_E , are computed using the method of Louis et al. (1982).

References

- Balmaseda, M. 2012. Initialization techniques in seasonal forecasts. In: *Proceedings of the ECMWF Seminar on Seasonal Prediction: Science and Applications*, ECMWF, Reading, UK, 3–7 September 2012, pp. 217–236.
- Balmaseda, M. and Anderson, D. 2009. Impact of initialization strategies and observations on seasonal forecast skill. *Geophys. Res. Lett.* **36**, L01701.
- Bannister, R. N. 2008a. A review of forecast error covariance statistics in atmospheric variational data assimilation. I: characteristics and measurements of forecast error covariances. *Q. J. Roy. Meteorol. Soc.* **134**, 1955–1970.
- Bannister, R. N. 2008b. A review of forecast error covariance statistics in atmospheric variational data assimilation. II: modelling the forecast error covariance statistics. *Q. J. Roy. Meteorol. Soc.* **134**, 1971–1996.
- Beljaars, A. C. M. 1995. The impact of some aspects of the boundary layer scheme in the ECMWF model. In: *Proceedings of the ECMWF Seminar on Parametrization of Sub-grid Scale Physical Processes*, ECMWF, Reading, UK, 5–9 September 1994, pp. 125–161.
- Beljaars, A. C. M. and Viterbo, P. 1998. The role of the boundary layer in a numerical weather prediction model. In: *Clear and Cloudy Boundary Layers* (eds. A. A. M. Holtslag and P. G. Duynkerke). Royal Netherlands Academy of Arts and Sciences, Amsterdam, pp. 287–304.
- Brassington, G. B., Martin, M. J., Tolman, H. L., Akella, S., Balmaseda, M. and co-authors. 2015. Progress and challenges in short- to medium-range coupled prediction. *J. Oper. Oceanogr.* (in press).
- Courtier, P. 1997. Dual formulation of four-dimensional variational assimilation. *Q. J. Roy. Meteorol. Soc.* **123**, 2449–2461.
- Courtier, P., Thépaut, J.-N. and Hollingsworth, A. 1994. A strategy for operational implementation of 4D-Var, using an incremental approach. *Q. J. Roy. Meteorol. Soc.* **120**, 1367–1387.
- Daley, R. 1991. *Atmospheric Data Analysis*. Cambridge University Press, Cambridge, UK.
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P. and co-authors. 2011. The ERA-interim reanalysis: configuration and performance of the data assimilation system. *Q. J. Roy. Meteorol. Soc.* **137**(656), 553–597.
- ECMWF IFS documentation. 2001–2013a. *Part III: Dynamics and Numerical Procedures*. European Centre for Medium-Range Weather Forecasts. Online at: <http://old.ecmwf.int/research/ifsdocs/>
- ECMWF IFS documentation. 2001–2013b. *Part IV: Physical Processes*. European Centre for Medium-Range Weather Forecasts. Online at: <http://old.ecmwf.int/research/ifsdocs/>
- Fowler, A., Bannister, R. and Eyre, J. 2012. A new floating model level scheme for the assimilation of boundary layer top inversions: the univariate assimilation of temperature. *Q. J. Roy. Meteorol. Soc.* **138**, 682–698.
- Gauthier, P., Charette, C., Fillion, L., Koclas, P. and Laroche, S. 1999. Implementation of a 3D variational data assimilation system at the Canadian Meteorological Centre. Part 1: the global analysis. *Atmos. Ocean*. **37**, 103–156.
- Gauthier, P., Tanguay, M., Laroche, S., Pellerin, S. and Morneau, J. 2007. Extension of 3DVAR to 4DVAR: implementation of 4DVAR at the Meteorological Service of Canada. *Mon. Weather Rev.* **135**, 2339–2354.
- Huang, X.-Y., Xiao, Q., Barker, D. M., Zhang, X., Michalakes, J. and co-authors. 2009. Four-dimensional variational data assimilation for WRF: formulation and preliminary results. *Mon. Weather Rev.* **137**, 299–314.
- Janisková, M., Thépaut, J.-N. and Geleyn, J.-F. 1999. Simplified and regular physical parameterizations for incremental four-dimensional variational assimilation. *Mon. Weather Rev.* **127**, 26–45.
- Keenlyside, N. S., Latif, M., Jungclaus, J., Kornbluh, L. and Roeckner, E. 2008. Advancing decadal-scale climate prediction in the North Atlantic sector. *Nature*. **453**, 84–88.
- Laloyaux, P., Balmaseda, M., Dee, D., Mogensen, K. and Janssen, P. 2015. A coupled data assimilation system for climate reanalysis. *Q. J. Roy. Meteorol. Soc.* DOI: 10.1002/qj.2629.
- Laloyaux, P., Balmaseda, M., Mogensen, K., Janssen, P. and Dee, D. 2014. The ECMWF coupled assimilation system. In: *Proceedings of the 2014 Conference on Big Data from Space (BiDS'14)* (eds. P. Soille and P. G. Marchetti). ESA-ESRIN, Frascati, Italy, 12–14 November 2014, pp. 16–19. DOI: 10.2788/1823.
- Large, W. G., McWilliams, J. C. and Doney, S. C. 1994. Oceanic vertical mixing: a review and a model with non-local boundary layer parameterization. *Rev. Geophys.* **32**, 363–403.
- Laroche, S., Tanguay, M. and Delage, Y. 2002. Linearization of a simplified planetary boundary layer parameterization. *Mon. Weather Rev.* **130**, 2074–2087.
- Lawless, A. S. 2012. *International Workshop on Coupled Data Assimilation*. Final report. University of Reading, UK. Online at: <http://www.esa-da.org/content/d1-report-international-workshop-coupled-data-assimilation>
- Lawless, A. S. 2013. Variational data assimilation for very large environmental problems. In: *Large Scale Inverse Problems: Computational Methods and Applications in the Earth Sciences* (eds. M. J. P. Cullen, M. A. Freitag, S. Kindermann, and R. Scheichl). Radon Series on Computational and Applied Mathematics 13. De Gruyter, Berlin/Boston, pp. 55–90.
- Lawless, A. S., Gratton, S. and Nichols, N. K. 2005. An investigation of incremental 4D-Var using non-tangent linear models. *Q. J. Roy. Meteorol. Soc.* **131**, 459–476.
- Lawless, A. S. and Nichols, N. K. 2006. Inner loop stopping criteria for incremental four-dimensional variational data assimilation. *Mon. Weather Rev.* **134**, 3425–3435.
- Lea, D. J., Mirouze, I., Martin, M. J., King, R. R., Hines, A. and co-authors. 2015. Using a new coupled data assimilation system to assess the HadGEM3 coupled model. *Mon. Weather Rev.* (in review).
- Lellouche, J.-M., Le Galloudec, O., Drévilion, M., Régnier, C., Greiner, E. and co-authors. 2013. Evaluation of global monitoring and forecasting systems at Mercator Océan. *Ocean Sci.* **9**, 57–81.

- Louis, J. F. 1979. A parametric model of vertical eddy fluxes in the atmosphere. *Bound. Layer Meteorol.* **17**, 187–202.
- Louis, J. F., Tiedtke, M. and Geleyn, J. F. 1982. A short history of the operational PBL parameterization at ECMWF. In: *ECMWF Workshop on Boundary Layer Parametrization*, ECMWF, Reading, UK, 25–27 November 1982, pp. 59–80.
- Mahfouf, J.-F. 1999. Influence of physical processes on the tangent-linear approximation. *Tellus A.* **51**, 147–166.
- Meehl, G. A., Goddard, L., Boer, G., Burgman, R., Branstator, G. and co-authors. 2014. Decadal climate prediction: an update from the trenches. *Bull. Am. Meteorol. Soc.* **95**, 243–267.
- Murphy, J., Kattsov, V., Keenlyside, N., Kimoto, M., Meehl, G. and co-authors. 2010. Towards prediction of decadal climate variability and change. *Proc. Environ. Sci.* **1**, 287–304.
- Navon, I. M., Zou, X., Derber, J. and Sela, J. 1992. Variational data assimilation with an adiabatic version of the NMC spectral model. *Mon. Weather Rev.* **120**, 1433–1466.
- Pierce, D. W., Barnett, T. P., Tokmakian, R., Semtner, A., Maltrud, M. and co-authors. 2004. The ACPI project, element 1: initializing a coupled climate model from observed conditions. *Clim. Change*. **62**, 13–28.
- Pohlmann, H., Jungclaus, J. H., Köhl, A., Stammer, D. and Marotzke, J. 2009. Initializing decadal climate predictions with the GECCO oceanic synthesis: effects on the North Atlantic. *J. Clim.* **22**, 3926–3938.
- Rabier, F., Jarvinen, H., Klinker, E., Mahfouf, J. F. and Simmons, A. 2000. The ECMWF operational implementation of four-dimensional variational assimilation. I: experimental results with simplified physics. *Q. J. Roy. Meteorol. Soc.* **126**, 1143–1170.
- Rawlins, F., Ballard, S. P., Bovis, K. J., Clayton, A. M., Li, D. and co-authors. 2007. The Met Office global four-dimensional variational data assimilation scheme. *Q. J. Roy. Meteorol. Soc.* **133**, 347–362.
- Ritchie, H., Temperton, C., Simmons, A., Hortal, M., Davies, T. and co-authors. 1995. Implementation of the Semi-Lagrangian method in a high-resolution version of the ECMWF forecast model. *Mon. Weather Rev.* **123**, 489–514.
- Shanno, D. F. 1978. On the convergence of a new conjugate gradient algorithm. *SIAM J. Numer. Anal.* **15**, 1247–1257.
- Shanno, D. F. and Phua, K. H. 1980. Remark on algorithm 500 – a variable method subroutine for unconstrained nonlinear minimization. *ACM Trans. Math. Software.* **6**, 618–622.
- Simmons, A. J. and Burridge, D. M. 1981. An energy and angular-momentum conserving vertical finite-difference scheme and hybrid vertical coordinates. *Mon. Weather Rev.* **109**, 758–766.
- Stewart, R. H. 2008. *Introduction to Physical Oceanography*. Department of Oceanography, Texas A&M University. Online at: <http://oceanworld.tamu.edu/index.html>
- Sugiura, N., Awaji, T., Masuda, S., Mochizuki, T., Toyoda, T. and co-authors. 2008. Development of a four-dimensional variational coupled data assimilation system for enhanced analysis and prediction of seasonal to interannual climate variations. *J. Geophys. Res.* **113**, C10017.
- Takaya, Y., Vitart, F., Balsamo, G., Balmaseda, M., Leutbecher, M. and co-authors. 2010. *Implementation of an Ocean Mixed Layer Model in IFS*. Technical Memorandum No. 622, European Centre for Medium-Range Weather Forecasts, Reading, UK.
- Tardif, R., Hakim, G. J. and Snyder, C. 2014. Coupled atmosphere-ocean data assimilation experiments with a low-order climate model. *Clim. Dyn.* **43**, 1631–1643.
- Thépaut, J.-N., Courtier, P., Belaud, G. and Lemaître, G. 1996. Dynamical structure functions in a four-dimensional variational assimilation: a case study. *Q. J. Roy. Meteorol. Soc.* **122**, 535–561.
- Thépaut, J.-N., Hoffman, R. N. and Courtier, P. 1993. Interactions of dynamics and observations in a four-dimensional variational assimilation. *Mon. Weather Rev.* **121**, 3393–3414.
- Viterbo, P., Beljaars, A., Mahfouf, J.-F. and Teixeira, J. 1999. The representation of soil moisture freezing and its impact on the stable boundary layer. *Q. J. Roy. Meteorol. Soc.* **125**, 2401–2426.
- Weaver, A. T., Vialard, J. and Anderson, D. L. T. 2003. Three- and four-dimensional variational assimilation with a general circulation model of the tropical pacific ocean. Part I: formulation, internal diagnostics, and consistency checks. *Mon. Weather Rev.* **131**, 1360–1378.
- Woolnough, S., Vitart, F. and Balmaseda, M. A. 2007. The role of the ocean in the Madden-Julian oscillation: implications for MJO prediction. *Q. J. Roy. Meteorol. Soc.* **133**, 117–128.
- Zhang, S., Harrison, M. J., Rosati, A. and Wittenberg, A. 2007. System design and evaluation of coupled ensemble data assimilation for global oceanic climate studies. *Mon. Weather Rev.* **135**, 3541–3564.